What Drives Risky Prescription Opioid Use?
Evidence from Migration

Amy Finkelstein
MIT and NBER

Matthew Gentzkow
Stanford and NBER

Dean Li
MIT[*]

October 2024

Abstract

We develop and estimate a dynamic model of risky prescription opioid use that allows us to unpack the role of person- and place-specific drivers of the opioid epidemic and to assess the impact of state opioid policies. Event studies indicate that, among adults receiving federal disability insurance from 2006 to 2019, moves to states with higher rates of risky use produce an immediate jump in the probability of risky use, followed by an additional gradual increase for the next several years. Using a potential outcomes framework, we show how these results map to the person- and place-specific factors in the model. Model estimates imply large effects of place on both the likelihood of transitioning to addiction and the availability of prescription opioids; they also indicate that these place effects change significantly when state laws restricting pain clinics are enacted. A one standard deviation reduction in all place effects would have reduced risky use by about 40 percent over our study period. One particular source of place effects, pain clinic laws, reduced risky use by 5 percent, but could have reduced it by 30 percent if they had been enacted earlier, with much of this magnification operating through the dynamics of addiction.

Keywords: Opioids; opioid epidemic; disability; geographic variation
JEL: H51, I12

[*]We thank the National Institute on Aging (Finkelstein R01-AG032449), the Social Security Administration (Finkelstein RDR18000003), and the Stanford Institute for Economic Policy Research (SIEPR) for financial support. We are grateful to Jonathan Gruber, Ellen Meara, David Molitor, Molly Schnell, and seminar participants at the 2021 NBER Summer Institute on Health Economics and the 2022 ASSA Session on the Economics of the Ongoing Opioid Crisis for helpful discussions, and to Alden Cheng, Paul Friedrich, Ken Jung, Geoffrey Kocks, Abby Ostriker, Wesley Price, Anna Russo, Steven Shi, and Yuci Zhou for excellent research assistance. We are especially grateful to Heidi Williams, who played a key role in the genesis of this project and made important contributions throughout its early stages. The research reported herein was performed pursuant to grant RDR18000003 from the US Social Security Administration (SSA) funded as part of the Retirement and Disability Research Consortium and is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No 214106. The opinions and conclusions expressed are solely those of the author(s) and do not represent the opinions or policy of SSA, any agency of the Federal Government, or the NBER. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States Government or any agency thereof.

1 Introduction

The opioid epidemic is one of the most important US public health crises of recent decades. In 2017, deaths from opioids were more than double the number of deaths from homicide, and an order of magnitude higher than cocaine-related deaths at the height of the 1980s crack epidemic (United States General Accounting Office 1991; Kochanek et al. 2019; National Center for Health Statistics 2021). A distinguishing feature of the origin of the epidemic is the role played by legal prescription opioids such as oxycodone (OxyContin) and hydrocodone (Vicodin). The epidemic is generally viewed as having arisen from physicians' increased willingness to prescribe legal opioids, which in turn was linked to pharmaceutical firms' marketing efforts and the medical profession's decision to recognize pain as a "fifth vital sign" (Quinones 2016; Case and Deaton 2020; Alpert et al. 2022; Maclean et al. 2022). But there has also been wide geographic variation in the evolution and the intensity of the opioid epidemic, with prescribing per capita varying by a factor of more than two between the 10th and 90th percentile states (McDonald, Carlson, and Izrael 2012; Schieber et al. 2019), and Appalachia, New England, and the West (Guy et al. 2017; Katz 2017) being particularly hard hit.

In this paper, we estimate a dynamic model of risky prescription opioid use ("risky use") that allows us to unpack the drivers of the prescription opioid epidemic, and to better understand the impact of opioid-related policies. We decompose drivers along two dimensions. First, we separate the roles of person-specific and place-specific factors. Person-specific factors include characteristics such as age, smoking, childhood adversity, mental health status, and prior substance use, whose importance has been emphasized in the medical literature (Ives et al. 2006; Sullivan et al. 2010; Fischer et al. 2012; Webster 2017). They also include individual human capital, which may be a contributor to what Case and Deaton (2015; 2017; 2020) term "deaths of despair." Place-specific factors include the propensity of local physicians to prescribe opioids for what they believe to be legitimate reasons (Barnett, Olenski, and Jena 2017; Schnell and Currie 2018; Eichmeyer and Zhang 2022), the availability of unscrupulous providers and "pill mill" pain clinics (Rutkow et al. 2015; Lyapustina et al. 2016), and the presence of policies such as triplicate prescription programs (Alpert et al. 2022), naloxone access laws (Doleac and Mukherjee 2018; Rees et al. 2019), laws regulating pain clinics ("pill mill laws") (Meara et al. 2016; Kaestner and Ziedan 2023), initial opioid prescribing limits for opioid-naive patients[1] ("prescribing limits") (Sacks et al. 2021), and prescription drug monitoring programs ("PDMPs") (Kilby 2015; Meara et al. 2016; Buchmueller and Carey 2018; Kaestner and Ziedan

[1]Opioid-naive patients are defined in these policies and the medical literature as patients who do not have recent prior opioid prescriptions.

2023).[2]

Second, we distinguish between place factors that impact the likelihood that individuals transition into and out of addiction and place factors that influence the ease with which addicted individuals can engage in risky use. We refer to the former as the *addiction* channel and the latter as the *availability* channel. The addiction channel is affected by factors like the willingness of physicians to give first-time prescriptions to patients suffering from pain. The availability channel is affected by factors like the presence of unscrupulous providers or the difficulty of receiving multiple or overlapping prescriptions.

In our model, an individual may transition into or out of addiction with probabilities that depend on both person- and place-specific factors. Whether or not she engages in risky use depends on her addiction state as well as place-specific availability factors. We allow both place-specific addiction transitions and place-specific availability to be impacted by policies. The model yields rich predictions for variation in the dynamics of the epidemic across space and time. It suggests that changes in risky use when people move across areas can identify the relative importance of person- and place-specific factors, while the time pattern of those changes helps distinguish between the two different types of place factors. The impact of policies is identified by the way the dynamics around moves change following policy implementation.

We use data on prescription opioid use from 2006 to 2019 among adults enrolled in the Social Security Disability Insurance (SSDI) program, the federal public insurance program for the disabled. Opioid use is especially prevalent in this population; roughly half of SSDI recipients received an opioid prescription each year during our sample period (Morden et al. 2014; Meara et al. 2016), a rate considerably higher than the rest of the population. We take advantage of rich panel data on prescription drug use for about 3 million SSDI recipients. The fact that the SSDI population has a relatively fixed level of government benefits and extremely tight limits on additional earnings lets us rule out large changes in individual income or employment as drivers of behavior change around moves. At the same time, it may also attenuate the variation in economic circumstances whose role in the opioid epidemic has been highlighted by Case and Deaton (2015; 2017; 2020).

To define risky prescription opioid use, we follow a standard approach in the literature (e.g., Larrick 2014; Meara et al. 2016; Ciesielski et al. 2017) and measure risky use for each individual-year with an indicator for prescriptions that result in an average daily morphine-equivalent dose (MED) of more than 120 mg in any calendar quarter. This measure has previously been found to be strongly correlated with adverse outcomes such as opioid poisonings and opioid dependence (Sullivan et al. 2010; Bohnert et al. 2011), and

[2]Area-specific peer effects (Khan et al. 2019; Powell, Pacula, and Taylor 2020) will also be captured as place factors in our framework.

we confirm this is true in our sample as well.

We begin with two pieces of reduced form analysis. First, we exploit staggered policy introductions across states to analyze the impact of three of the most prominent categories of state opioid policies: pill mill laws, PDMPs, and prescribing limits. Prescribing limits were particularly aimed at limiting transitions of new users into addiction, while pill mill laws and PDMPs were motivated by concerns about inappropriate prescriptions to those who are already addicted. However, recent work has noted that all three policy categories have the potential to affect both the addiction and availability channels.[3] Our results suggest that pill mill laws have the strongest impacts on risky use in our sample; we therefore incorporate pill mill laws directly in our model.

Second, we present event study analyses of changes in risky use when individuals move across states. We show that individuals' probability of risky use increases immediately when they move to states with higher rates of risky use, and falls immediately when they move to states with lower rates of risky use. These effects then increase in the years post-move. Prior opioid users have a larger on-impact effect, whereas patients without prior prescriptions have no on-impact change but a steady increase post-move. Using a potential outcomes framework, we show formally how the event study coefficients map back to the person- and place-specific factors in our model.

We then estimate the model by Generalized Method of Moments, using pre-move and post-move moments from about 16,600 separate mover cohorts that capture the same variation explored in our event studies, as well as moments based on non-movers which are used to identify the impacts of pill mill laws. We define a mover cohort by its origin and destination state, and by the calendar year in which the move occurs. The empirical implementation of our model allows for the possibility that movers are non-randomly selected, with person-specific factors that may be correlated with those of non-movers in both their origin and their destination; the key identifying assumption is that the pattern of selection does not change systematically with the time relative to the move, and in particular that it does not change discretely on move. We also allow for the possibility that pill mill policies are phased in gradually, with an implementation path we estimate from data. Predicted moments from the model align well with their empirical counterparts, and simulations based on the estimated parameters closely match the estimated mover event studies as well as estimates of the impact of pill mill regulations based on their staggered introduction across states, even though neither of these reduced form findings were explicitly targeted in estimation.

[3]For example, Alpert, Dykstra, and Jacobson (2024a) find that PDMPs can affect both prior users and opioid-naive patients through introducing new physician hassle costs to opioid prescribing. Likewise, the literature on pill mill policies suggests that they may have affected both risky opioid use and potentially appropriate use (Rutkow et al. 2015; Lyapustina et al. 2016; Kaestner and Ziedan 2023).

The model estimates reveal substantial geographic heterogeneity in the drivers of the opioid epidemic, with large variation in both person-based and place-based factors. For example, we find that availability played a larger relative role in the West and Northeast, while place-specific and person-specific addiction factors were more important in the South and Appalachia. We also find that pill mill policies reduce both availability and addiction parameters, with their largest effects on probabilities of transitioning into addiction. Our estimated model parameters are correlated in intuitive ways with various area-level characteristics. For example, the average prescription opioid dose given after surgery is strongly positively correlated with place-based addiction factors, which is consistent with finding that variation in initial opioid prescriptions have strong effects on downstream addiction outcomes (Eichmeyer and Zhang 2022). Triplicate prescribing laws are strongly negatively correlated with place-based addiction factors, which is consistent with the large effects of triplicate-prescribing programs found by Alpert et al. (2022) operating through lowering the likelihood of individuals transitioning into addiction.

We present two main sets of counterfactuals from the estimated model. First, we decompose the impact of different channels in contributing to risky prescription opioid use over our 2006 to 2019 sample period. Consistent with evidence from the mover event studies, we estimate substantial roles for both person-specific and place-specific factors. A one standard deviation reduction in all person-specific factors would have reduced the average rate of risky use during this time period by 33 percent. Of course, these person-specific factors in turn may have been shaped by the long-run impacts of place-specific factors; in that sense, one can think of our estimates of the role of place-specific factors as a lower bound on their impact. A one standard deviation reduction in all place-specific factors would have reduced the rate of risky use by 41 percent. Reducing place-specific availability has a relatively larger effect in the very short-run than reducing place-specific addiction transitions, but over the full fourteen years of our study period the impact of the latter is 2.5 times larger than the impact of the former.

Second, we consider the impact of pill mill laws, which our model estimates indicate affect both place-based addiction and availability channels. We estimate that these laws decreased risky use by 4.4 percent in the eleven implementing states during our sample period. We also show that if these laws had been in place in these states since the first wave of OxyContin-related marketing in 1996, their effects would have been substantially larger, decreasing risky use by 30 percent over the same sample period. The effect of early implementation is magnified because the policy effects compound through the dynamics of addiction.

An important limitation is that we measure risky use of prescription opioids, and do not measure illegal opioid use. If factors which we estimate reduce risky use of prescription opioids are associated with substitution toward illegal opioids such as heroin or fentanyl, then the welfare effects of reductions in risky

prescription opioids could be attenuated or even reversed. To shed some light on this issue, we look at how our estimated availability and addiction parameters – which drive risky use of prescription opioids – correlate with rates of adverse health outcomes such as opioid poisonings, which capture impacts of both prescription opioids and illegal opioids. We find that our availability and addiction parameters are strongly correlated with higher rates of adverse outcomes, suggesting that any substitution is unlikely to fully offset health benefits from reductions in risky use of prescription opioids.

Our analysis relates to a large literature on the causes and consequences of the opioid epidemic. We connect to work that separately examines the importance of person and place factors in driving the epidemic. Much of this literature focuses on correlations between measures of the epidemic and changes in other factors. Some researchers, including Case and Deaton (2015; 2017; 2020), have pointed to the important role of individuals' economic circumstances and related demand-side factors (Carpenter, McClellan, and Rees 2017; Hollingsworth, Ruhm, and Simon 2017; Dean and Kimmel 2019; Venkataramani et al. 2020). Others have argued that the patterns of increased drug deaths and measures of risky opioid use across states suggest an important role for the availability of opioids, and little, if any, role for economic decline (Ruhm 2018; Currie, Jin, and Schnell 2019; Currie and Schwandt 2021; Cutler and Glaeser 2021).[4] Still others, such as Charles, Hurst, and Schwartz (2019), conclude that both economic conditions and opioid supply have played a role in local opioid use and deaths. To complement this work, we develop an explicit model of opioid addiction, and we use an empirical strategy based on individuals moving across states and the staggered timing of pill mill laws, to estimate the model and decompose causal drivers.

We also contribute to a literature that analyzes the effects of opioid-related policies. We build on several national analyses of PDMPs, prescribing limits, and pill mill laws (Meara et al. 2016; Buchmueller and Carey 2018; Sacks et al. 2021; Kaestner and Ziedan 2023) by re-analyzing these policies within our sample. We further provide a framework for estimating the effects of these policies and understanding their dynamics within a model of opioid addiction.

More broadly, our approach fits into a large and growing literature using changes in location to separate the effects of individual characteristics from the effects of geographical or institutional factors on a wide range of outcomes (e.g. Card, Heining, and Kline 2013; Chetty, Friedman, and Rockoff 2014; Chetty and Hendren 2018a,b; Finkelstein, Gentzkow, and Williams 2016; Song et al. 2010; Molitor 2018).[5] Most of this prior literature has assumed separability between person-specific and place-specific factors in the outcome

[4] Other work suggests that economic decline may be a result of risky use itself (Laird and Nielsen 2016; Krueger 2017; Aliprantis, Fee, and Schweitzer 2023; Harris et al. 2020).

[5] Most closely related to the current paper, Laird and Nielsen (2016) exploit individual moves across municipalities in Denmark to generate quasi-random matches between individuals and physicians and estimate that treatment by a physician with a higher rate of opioid prescribing is associated with both more use of prescription opioids and a decline in labor market activity.

equation.[6] We extend the approach to allow for non-separability in the outcome equation of person-specific and place-specific factors, as well as for distinct place-specific channels. We also extend prior work by allowing for a specific form of selection of movers tailored to our setting.

Finally, our analysis of geographic variation in risky prescription opioid use relates to the large literature on the causes and consequences of geographic variation in health care utilization and treatment. Skinner (2011) provides a review of this literature.

The rest of the paper proceeds as follows. Section 2 describes the setting and data and presents the reduced form evidence of the impact of pill mill laws. Section 3 introduces the model. Section 4 presents estimates from mover event studies and discusses how the resulting patterns can be interpreted in light of the model. Section 5 discusses the parameterization of the model and its estimation. Section 6 presents the model estimates and our analysis of the impacts of counterfactual place-based policies. The last section concludes.

2 Setting, Data, and Policies

2.1 Setting

We study risky prescription opioid use over the fourteen-year period from 2006 to 2019. This period spans the initial surge of overdose deaths from prescription opioids, the subsequent rise in heroin-related mortality, and the more recent increase in synthetic opioid overdoses. Prescription opioids play a large role throughout this entire period (see Appendix Figure A.1).

Our baseline analysis focuses on prescription opioid use among a particularly high-risk population: adults enrolled in the Social Security Disability Insurance (SSDI) program, the federal public insurance program for the disabled. SSDI provides both income and medical insurance (Medicare) to workers experiencing long-term disability. Medical insurance coverage begins two years after eligibility for income payments. In 2010, over eight million disabled workers (and two million spouses and dependents) received SSDI benefits, and SSDI expenditures comprised over seven percent of federal non-defense spending; about one-third of these SSDI expenditures reflect Medicare costs (Autor 2015). Once on SSDI, it is rare for individuals to exit SSDI and return to the labor force (Autor and Duggan 2006). While on SSDI, recipients' income and labor market status tend to be stable; they are not allowed to have monthly labor market earnings above a low limit, which was $1,090 per month in 2015 (Social Security Administration 2017).

We focus on the SSDI population for several reasons. First, the stringent limits on labor market ac-

[6] One exception is the study of habit formation in consumer purchases by Bronnenberg, Dubé, and Gentzkow (2012).

tivity in the SSDI program minimize the likelihood of large changes in individual employment or income around the time of moves, allowing us to isolate on-impact place effects that are unlikely to be driven by changing economic circumstances. Second, we are able to leverage detailed, individual-level panel data on prescription opioid use in this population. And third, rates of prescription opioid use are high among this population (Morden et al. 2014; Meara et al. 2016). While the first two reasons also apply to those who qualify for Medicare because they are 65+, the last one does not; although SSDI recipients account for approximately a fifth of Medicare enrollees, they account for nearly four-fifths of opioid overdose deaths among the Medicare population (Kuo, Raji, and Goodwin 2019). However, to explore the robustness of our analyses to alternative populations, we construct an alternative sample of elderly Medicare (65+) enrollees. We describe the construction of this sample and present a parallel set of analyses in Appendix C.

2.2 Data

We use data from SSDI recipients in a 20 percent random sample of Medicare beneficiaries from 2006 through 2019. We focus on the approximately three-quarters of SSDI recipients who are also enrolled in Medicare Part D, the voluntary, heavily subsidized prescription drug benefit program that has been available to Medicare enrollees since 2006; the Part D enrollment rate among SSDI Medicare recipients is slightly higher than the 60 percent rate among elderly (65+) Medicare recipients (Cubanski, Neuman, and Damico 2016). Medicare recipients can enroll in Part D through a stand-alone Part D plan or through Medicare Advantage—a set of private insurance plans that offer an alternative form of health insurance to traditional Medicare; we observe their prescription drug fills in either case. In addition, adults with sufficiently low earnings prior to the onset of disability also qualify for Medicaid—the public health insurance program for qualifying low income adults—and this dual eligibility for Medicaid and Medicare also ensures that beneficiaries have Part D coverage.

We are able to follow Part D enrollees over time in an individual-level panel, and to observe basic demographic information including gender, age, race, Medicaid enrollment, and zip code of residence, which we define annually as the address on file for Social Security payments.[7] We use the detailed, claim-level data to observe the drug and dosage of filled prescriptions; prior work has also used the Medicare Part D prescription drug claims data to measure prescription opioid use (e.g. Morden et al. 2014; Meara et al. 2016; Buchmueller and Carey 2018). For the approximately three-quarters of our enrollee-years who are not

[7] We obtain data on zip code of residence from the Medicare Denominator File for years 2006-2008 and Beneficiary Summary File (BSF) for years 2009-2019. In the Medicare Denominator File, a beneficiary's zip code of residence each year is determined by her address on file as of March 31st of the following year; in the BSF, a beneficiary's zip code of residence is determined by her address on file as of December 31st that year. See https://www.resdac.org/articles/medicare-eligibility-and-enrollment-files-rif-versions/ for more information.

enrolled in Medicare Advantage, we also observe inpatient and outpatient claims which we use to develop various measures of adverse opioid outcomes.

Sample Construction From our original 20 percent sample of all Medicare enrollees (18.8 million enrollees; 153 million enrollee-years) we limit ourselves to the approximately one-fifth of enrollees whose original source of Medicare eligibility was from SSDI, and then further restrict to the approximately two-thirds of enrollee-years with 12 months of Medicare Part D coverage. We exclude from this sample person-years in which the enrollee is older than 99. This produces a sample of about 3.1 million enrollees (22.9 million enrollee-years).

Our baseline geographic unit of analysis is a state, although we also present results for other levels of geography such as county or commuting zone. We define individuals to be "non-movers" if their state of residence is the same throughout our sample period. We define individuals to be "movers" if their state of residence changes exactly once during this period.

Starting from a sample of 320,000 enrollees who change their state of residence at least once, we impose several additional restrictions to arrive at our final sample of 99,729 movers. We exclude the approximately 24,000 movers who moved during the last year in the sample, since we cannot observe their post-move behavior. We exclude about 110,000 movers who changed their state of residence more than once during our study period. We limit analysis to at most five years pre-move, the move year, and five years post-move, thus excluding about 370,000 enrollee-years outside this window. We follow the approach developed in Finkelstein, Gentzkow, and Williams (2016) and exclude about one-quarter of the remaining movers whose share of prescription claims in their destination state, among prescription claims (for any drug) in either their origin or destination state, is not at least 0.75 higher in the post-move years relative to the pre-move years.[8] Finally, we exclude approximately 40,000 movers who we do not observe in the year before their move.[9]

[8] Appendix Figure A.2 shows that our approach successfully identifies the timing of moves, with about 50 percent of origin or destination claims in the move year located in either a mover's origin or her destination. Note that we do not directly observe the location of Part D claims, but we do observe the identification number of the prescriber. We define the prescriber's location in a given year as the state where they have at least 60 percent of their inpatient, outpatient, and carrier claims for that year. We do not define a location for any prescriber that does not have at least 60 percent of their annual claims within a single state. On average, for our non-mover SSDI population, we estimate that about 97 percent of opioid prescriptions filled are prescribed by a doctor who practices within the individual's state of residence. Note that our measure of the change in the claim share in a given location is not defined for movers who do not have at least one claim both pre- and post-move. We exclude these cases if: (i) they have no post-move prescription claims and a pre-move destination prescription claim share greater than 0.05; (ii) they have no pre-move prescription claims and a post-move destination prescription claim share less than 0.95. See Finkelstein, Gentzkow, and Williams (2016) for a more detailed discussion.

[9] In order to keep the sample consistent throughout our event study analysis and model estimation, we further restrict our sample to movers who have a matched non-mover observation in the same year, matching on origin state, five-year age bin, race, and sex. This excludes only an additional 448 enrollee-years, as nearly all mover-years are matched.

Measuring Risky Prescription Opioid Use Opioids are both a risky addictive drug and a critical source of relief for patients suffering acute pain. This makes it difficult to determine with certainty which prescriptions are consumed or diverted for non-medical purposes, and which are part of a medically appropriate treatment plan. Even in a clinical setting, physicians may struggle to identify misuse (Parente et al. 2004).

While there is no consensus measure of misuse among clinicians or medical researchers (Turk, Swanson, and Gatchel 2008; Sullivan et al. 2010), the medical literature studying the opioid epidemic has developed several proxies for risky prescription opioid use based on prescription data (Hall et al. 2008; White et al. 2009; Sullivan et al. 2010; Cepeda et al. 2012; Rice et al. 2012; Logan et al. 2013; Jena et al. 2014; Larrick 2014; Morden et al. 2014; Meara et al. 2016). These measures identify patterns in prescriptions at the person-year level that are correlated with adverse drug outcomes such as opioid dependence, emergency room visits, and overdose deaths (Braden et al. 2010; Dunn et al. 2010; Bohnert et al. 2011; Logan et al. 2013; Jena et al. 2014; Ciesielski et al. 2017; Brat et al. 2018; Klimas et al. 2019; Rough et al. 2019; Wei et al. 2019).

Our primary measure of risky use is an indicator for an individual filling prescriptions that result in an average daily morphine equivalent dosage (MED) of more than 120 mg in any calendar quarter of the year. This dosage is above the 96th percentile of enrollee-quarter observations in our data (and above the 93rd percentile among enrollees who are prescribed any opioids in the year); it is about six times higher than the average quarterly dosage in our sample, and three and a half times higher than the average quarterly dosage among those with any opioid use in the year. Using the typical Vicodin dosage, it would correspond to 24 Vicodin pills each day during a three-month period. Our use of this measure follows Sullivan et al. (2010), Larrick (2014), Meara et al. (2016), and Ciesielski et al. (2017).

We also show robustness of our mover event studies to two other measures of risky use used in the literature: an indicator for filling prescriptions with four or more unique physicians in a year, and an indicator for filling a new opioid prescription before the end of a previous prescription during the year. Appendix A provides more detail on the construction of each measure and their correlations with each other and with subsequent adverse outcomes in our data. We selected our primary measure of risky use based on the observation that it is the most predictive of subsequent adverse opioid outcomes (i.e., opioid poisonings or opioid use disorders) in our data (see Appendix Table A.1).

Finally, for some mover analyses, it will be useful to distinguish among patients based on their prior opioid history. Following the standard definitions in the literature (Sun et al. 2016; Deyo et al. 2017), we define a mover as "opioid naive" if she filled no opioid prescriptions in the year before the move, and as a "prior user" if she filled an opioid prescription in the year before the move.

Summary Statistics Table 1 reports summary statistics for our study population of movers (column 1). Slightly over one-half are female, and approximately three-quarters are white. Slightly more than one-half receive Medicaid, and their average age is 58. In just under half of enrollee-years there is at least one opioid prescription, which is consistent with the previously documented high rate of opioid use in the disabled population (Morden et al. 2014; Meara et al. 2016). About four percent of enrollee-years meet the definition for risky use. In the year prior to move, movers are slightly more likely to be opioid naives than prior users. Column 2 shows comparable statistics for our study population of non-movers. Relative to non-movers, movers are slightly more likely to be female, white, and younger. Movers also exhibit somewhat higher rates of prescription opioid use and of risky use.

Among our movers, we estimate that the average distance between their origins and destinations (measured between the population-weighted state centroids based on the 2010 census) is 797 miles, with a median move of 638 miles and a standard deviation of 617 miles. The median state receives 1,536 movers (about 2% of movers) and the mean state receives 1,955. Florida is the most common destination state (about 13% of movers), while California is the most common origin state (9.8% of movers). The least common destination is the District of Columbia (0.1% of movers) and the least common origin is North Dakota (0.2% of movers).

Figure 1 shows the distribution and time-path of the rates of risky prescription opioid use across states for our full sample of non-movers. Panel A presents a map of average risky use rates and highlights the regions that were most affected by the prescription opioid crisis. New England, Appalachia, the Southwest, and the Northwest are all particularly hard-hit. There is also considerable variation within regions. For example, within New England, the average risky use rate in New Hampshire is more than twice that of the neighboring state of Massachusetts, while in the Midwest, the risky use rate in Montana is more than two and a half times higher than that of the neighboring state of North Dakota.[10] These differences will prove useful for generating variation in place-based factors when individuals move across states. Panel B presents trends in risky prescription opioid use. Risky use grows in prevalence until 2013 and begins to fall afterwards, as the locus of the crisis shifted from prescription to illegal opioids. Notably, the spread between the 20th and 80th percentile risky use states remains stable over time, which is consistent with national changes in attitudes and guidelines playing similar roles across geographies.

[10] While we cannot directly compare the patterns of risky use as we measure it in our population to that of the general population, in Appendix C we show that national trends and state-level variation in opioid prescriptions per capita in our population are similar to that in the general population, although the level of prescriptions per capita is substantially higher for our disabled population.

2.3 State Opioid Policies and their Reduced Form Impacts

We consider three policies that were widely adopted during our time period and that have also been highlighted by the prior literature: PDMPs, pill mill laws, and prescribing limits.[11] Different states enacted these policies at different times, providing a key source of variation that we (and prior papers) use to identify their effects. Appendix B describes these policies, the timing of their implementation, and our reduced form estimates of their impacts in our non-mover study population. Appendix Figure A.3 presents the associated policy event studies, while Appendix Figure A.4 presents our estimates of each policy's average impact in its first five years, and, where feasible, compares our estimates to ones from the existing literature. For each specification, we estimate confidence intervals through 50 iterations of the Bayesian bootstrap procedure (Rubin 1981) which smooths bootstrap samples through reweighting rather than resampling observations. We find that pill mills are associated with a statistically significant 10% decline in risky use, but that PDMPs and prescribing limit policies are associated with substantially smaller and statistically insignificant effects on risky use. These findings are broadly consistent with existing estimates, many of which are imprecise.

Based on these findings, in our modeling below we focus our analysis of policy impacts on the effects of pill mill laws. We make two additional modeling choices based on our estimates as well as qualitative analyses of pill mill laws. First, our estimates suggest that the impact of the pill mill laws grows over time (see Appendix Figure A.3, Panel B). This could reflect a growing impact of a fully phased-in law, or a gradual phase-in of a law with a time-constant impact, or some combination of both. In practice, it appears that pill mill laws are gradually phased-in due to both logistical hurdles to implementation as well as subsequent modifications to the policy. Appendix Figure A.5 documents the process of subsequent modifications to initially-implemented pill mill laws, and Stone et al. (2020) provide a qualitative discussion of factors behind implementation delays based on interviews with policy-makers in several states that introduced pill mill laws.[12] As a result, our policy parameterization in the model estimation in Section 5 below will allow for a potential gradual phase-in of these laws.

Second, we allow these policies to affect both the availability and place-based addiction channels. Pill mill laws most directly target availability. They aim to reduce inappropriate prescribing by healthcare facilities that specialize in the diagnosis and management of chronic pain by defining a set of requirements

[11] PDMPs vary widely across states. A lot of the literature studying these policies focuses on the implementation of "must-access" mandates (Buchmueller and Carey 2018; Alpert, Dykstra, and Jacobson 2024a); these refer to a strengthened set of requirements which dictate that providers access the PDMP prior to prescribing or filling a prescription for a controlled substance. We follow this approach. Our analysis of PDMPs is likewise an analysis of must-access PDMPs.

[12] Potential sources of implementation delays provided by interviewees in Stone et al. (2020) include the time associated with mobilizing pill mill inspections, coordination with law enforcement, information delays associated with definitions and informing providers who "did not understand that they needed to register as a pain clinic to comply with the law", and the time required for external parties to raise flags about signals of potential non-compliance with the law.

and responsibilities for such facilities.[13] In practice, however, the policies may also affect transitions into and out of addiction. For one thing, the laws can apply quite broadly to healthcare facilities which aim to treat patients with chronic pain, many of whom are not already addicted to opioids.[14] Such facilities serve the sizable population suffering from chronic pain: estimates suggest that over 20% of U.S. adults experience chronic pain, and that nearly 7% of adults feel that this pain restricted their daily activities Rikard 2023. Furthermore, opioid prescribing behavior by physicians is highly concentrated among high-volume prescribers: in 2017, the top 1% of providers were responsible for 49% of all opioid doses and 27% of all opioid prescriptions Kiang et al. 2020. Thus, these policies may directly affect both the availability of large doses as well as the rate of first-time prescriptions and initiations into addiction more broadly.

3 Model

We build a dynamic model of opioid addiction and risky opioid use. We consider a population of people i living in locations j in years t, with $j(i,t)$ denoting i's location in year t. Each person is either a non-mover who stays in a single location in all periods or a mover who changes location exactly once. We denote the set of non-movers in location j by $\mathscr{I}_j$. We define a *cohort* of movers by their origin o, destination d, and move year m, and index these cohorts by c. The set of movers in cohort c is $\mathscr{I}_c$, and we let $m(c)$, $o(c)$, and $d(c)$ denote cohort c's move year, origin, and destination, respectively. For all $i \in \mathscr{I}_c$, we use $c(i)$ to refer to the cohort that they belong to. Finally, we use $\mathscr{C}_m$ to denote the set of all move cohorts.

In each year t, person i may be either addicted to prescription opioids or not addicted. We denote this addiction state by $a_{it} \in \{0,1\}$, with $a_{it} = 1$ indicating that i is addicted in year t. The addiction state evolves stochastically with time-varying probabilities of transitioning into or out of addiction that depend on both person- and place-specific factors. The probability of transitioning to addiction is given by $\pi^+_{jt} + \eta^+_i$ and the probability of transitioning out of addiction is given by $\pi^-_{jt} + \eta^-_i$, where the π and η terms represent place- and person-specific factors respectively.[15] Thus,

$$Pr(a_{i,t} = 1 | j(i,t) = j) = \begin{cases} 1 - \pi^-_{jt} - \eta^-_i & \text{if } a_{i,t-1} = 1 \\ \pi^+_{jt} + \eta^+_i & \text{if } a_{i,t-1} = 0. \end{cases}$$

[13] These responsibilities can include explicit certification requirements, requiring physician ownership, complying with specific prescribing restrictions, requiring physician examinations for patients, requiring drug-testing for high-risk patients, and keeping patient records.

[14] For example, Kentucky's definition of pain management facilities include all facilities that advertise for any type of pain management services (Kentucky Revised Statutes § 218A.175).

[15] We restrict the domain of these parameters to values where addiction transitions can be interpreted as probabilities (i.e., [0,1]). The domain restrictions applied in our estimation are discussed in Appendix E.

We let $\pi_{jt} \equiv \left[\begin{array}{cc} \pi_{jt}^{+} & \pi_{jt}^{-} \end{array} \right]$ and $\eta_i \equiv \left[\begin{array}{cc} \eta_i^{+} & \eta_i^{-} \end{array} \right]$.

The expected share $\overline{a}_{jt}$ of non-movers in location j who are addicted at time t is defined recursively as

$$\overline{a}_{jt} = \frac{1}{\left|\mathscr{I}_j\right|} \sum_{i \in \mathscr{I}_j} \left[\left(1 - \pi_{jt}^{-} - \eta_i^{-}\right) a_{i,t-1} + \left(\pi_{jt}^{+} + \eta_i^{+}\right)(1 - a_{i,t-1}) \right],$$

where we set $a_{i0} = 0$ and $\left|\mathscr{I}_j\right|$ denotes the number of non-movers in location j.

Our main observed outcome of interest is risky prescription opioid use. We treat this as binary, letting $y_{it} \in \{0,1\}$ be an indicator for whether person i engages in risky use in year t, and we let $\overline{y}_{ct}$ denote the mean of y_{it} among movers in cohort c. The probability a person who is addicted in a given year engages in risky use is given by γ_{jt}, a term that may depend on both place- and time-specific factors, including local supply-side conditions and opioid policies.[16] As a shorthand, we refer to the parameter γ_{jt} as "opioid availability" in location j and time t. Our baseline model assumes that non-addicted individuals never exhibit risky use; in Section 6.2 we show that our counterfactuals are robust to relaxing this restriction. The probability of risky use is thus given by

$$Pr(y_{it} = 1 | j(i,t) = j) = \begin{cases} \gamma_{jt} & \text{if } a_{it} = 1 \\ 0 & \text{if } a_{it} = 0. \end{cases}$$

Expected risky use conditional on the addiction state is $a_{it}\gamma_{jt}$.

In the empirical analysis, we will observe risky use y_{it} and the location of each individual, but not their underlying addiction state a_{it}. We cannot directly observe the addiction state a_{it}, since not all addicted individuals engage in risky use in a given year. The key identification challenge will therefore be to separately pin down the opioid availability parameters γ_{jt} and the place-specific and person-specific addiction parameters π_{jt} and η_i. As we will discuss, the patterns of risky use around moves to and from different locations will be crucial for identification.

A main focus of the analysis will be characterizing the component of risky use that can be attributed to the causal effect of place. For an individual who is addicted ($a_{it} = 1$), the causal effect of moving in period t from an original location $o(c(i))$ to a new location $d(c(i))$ on risky use in period t is $\left(\gamma_{d(c(i))t} - \gamma_{o(c(i))t}\right)$. The causal effect for a non-addicted individual is zero. Thus, the average causal effect of such a move in period t for a population in which share $\overline{a}$ is addicted is $\overline{a}\left(\gamma_{d(c(i))} - \gamma_{o(c(i))t}\right)$. Over time, the share addicted $\overline{a}$ will also be affected by the move via place-specific addiction parameters.

[16] In the empirical analysis below we allow for the possibility that movers are non-randomly selected from their origin population, which is why here we define γ_{jt} for a non-mover.

To define the causal effect of place in full generality, we can define a set of potential outcomes $y_{it}(\mathbf{h})$, where $\mathbf{h} = (j_1, j_2, ..., j_t)$ is a vector indicating the history of locations in which i lived in each year $1, ..., t$, and where $y_{it}(\mathbf{h})$ is the outcome y_{it} that would occur under history $\mathbf{h}$. We similarly define $\overline{a}_{ct}(\mathbf{h})$ to be the share of individuals in cohort c at time t who would be addicted were all of those individuals to have location history $\mathbf{h}$. We assume that a mover is in their origin location for $t < m$, is in their destination for $t > m$, and may be in either their origin or their destination in $t = m$. Note that our formulation implicitly rules out anticipatory effects of moving, as potential outcomes only depend on current and past locations.

We can then define the period-t average treatment effect on movers in cohort c of moving in year $m(c)$ relative to remaining in their origin as:

$$T_{ct} = E_{i \in \mathscr{I}_c}\left[y_{it}(\mathbf{h}_{ct}) - y_{it}\left(\mathbf{h}^0_{ct}\right)\right] \tag{1}$$

where $\mathbf{h}^0_{ct}$ is the t-history where all elements are equal to $o(c)$ and $\mathbf{h}_{ct}$ is the t-history in which the first $m(c)$ elements are $o(c)$ and the remaining elements are $d(c)$. It is straightforward to show that

$$T_{ct} = \begin{cases} 0 & t < m(c) \\ \overline{a}_{ct}\gamma_{d(c)t} - \overline{a}^0_{ct}\gamma_{o(c)t} & t > m(c) \end{cases} \tag{2}$$

where $\overline{a}_{ct} = \overline{a}_{ct}(\mathbf{h}_{ct})$ and $\overline{a}^0_{ct} = \overline{a}_{ct}\left(\mathbf{h}^0_{ct}\right)$. If movers were as-good-as-randomly selected from non-movers in their origin, $\overline{a}^0_{ct}$ would simply be equal to the share of non-movers who are addicted in the origin location $o(c)$ in period t. Our empirical implementation of the model, however, will allow for the possibility of non-random selection.

The period-t average treatment effect on movers in equation (2) consists of two distinct channels, reflecting the impact of place on both the share of addiction $(\overline{a}_{ct})$ and on the rate of risky use conditional on addiction. To see this more clearly, we can rewrite the term for $t > m(c)$ as

$$T_{ct} = \left(\overline{a}_{ct} - \overline{a}^0_{ct}\right)\gamma_{d(c)t} + \overline{a}^0_{ct}\left(\gamma_{d(c)t} - \gamma_{o(c)t}\right),$$

where the first term can be interpreted as the treatment effect due to the addiction channel and the second term can be interpreted as the treatment effect due to the availability channel.

4 Reduced Form Evidence from Moves

We present evidence of how the propensity for risky use changes with the timing and direction of moves, and how these patterns vary across individuals with different probabilities of addiction. We discuss how these patterns can be interpreted in light of our model.

4.1 Event Study Specification

We specify an event study regression that recovers an average of the cohort and period-specific average treatment effects defined in Section 3. We allow for fixed individual differences in risky use rates, arbitrary trends in those rates correlated with observables, and arbitrary trends in those rates around moves that are the same for all mover cohorts.[17] Our approach also follows recent work emphasizing the importance of allowing for heterogeneity in treatment effects across units treated at different times (e.g. de Chaisemartin and D'Haultfœuille 2020; Callaway and Sant'Anna 2021; Sun and Abraham 2021; Wooldridge 2021); in our context, this means that we allow for heterogeneous treatment effects based on mover cohort and observable covariates X_i.

To focus on event time, we let $r(c,t) \equiv t - m(c)$ index years relative to a cohort's move year. We use the subscript cr as shorthand for $c,t(c,r)$, where $t(c,r)$ is the calendar year corresponding to relative year r for cohort c. For example, $T_{cr} = T_{c,t(c,r)}$. We similarly use the subscript ir as shorthand for $i,t(c(i),r)$ when referring to movers. For example, for any mover i we have $y_{ir} = y_{i,t(c(i),r)}$. We define X_i to be a vector of observables for mover i consisting of the interaction between five-year age bin (as measured by birth year), race (White, Black, or other), and gender. Finally, for each mover cohort c, we define a sample of matched non-movers consisting of all non-movers who are observed in cohort c's pre-move year $r = -1$ in the cohort's origin $o(c)$.

We let $\hat{m}^X_{cr}$ denote the average of y_{ir} among the subset of c's matched non-movers with characteristics X. We then define $\hat{T}^X_{cr}$ to be the average of the difference in differences $(y_{ir} - y_{i,-1}) - \left(\hat{m}^X_{cr} - \hat{m}^X_{c,-1}\right)$ among i belonging to the set $\mathscr{I}^X_c$ of movers in cohort c with characteristics X. We let $\hat{T}_{cr}$ denote the average of $\hat{T}^X_{cr}$ across X, weighting by the number of movers in c with each set of characteristics X. Defining $\hat{T}^X_{cr}$ as a difference in difference adjusts for any fixed individual differences in propensity for risky use, and defining it conditional on X before averaging across X to obtain $\hat{T}_{cr}$ allows for arbitrary calendar year trends in y that are specific to each value of X. We show below that under an appropriate parallel trends assumption, $\hat{T}_{cr}$ is a consistent estimator for the treatment effect T_{cr}.

[17] These adjustments correct for selection of movers in a way that is flexible, although not directly tied to our model. When we parameterize our model in Section 5 below, we will correct for selection in a way that has a precise structural interpretation.

Finally, in the spirit of Finkelstein, Gentzkow, and Williams (2016), we scale these estimated treatment effects by a measure of the gap in risky use between a mover's origin and destination, allowing us to interpret the event study coefficients as the share of geographic differences closed by moving. Specifically, we define δ_{cr} to be the difference $\bar{y}_{d(c)r} - \bar{y}_{o(c)r}$ in the average rates of risky use among matched non-movers in period $t(c,r)$ in the cohort's destination $d(c)$ and origin $o(c)$.[18] To simplify the interpretation of our estimates in Section 4.2, we normalize δ_{cr} by subtracting the mean of this difference across all movers.[19] We let $\hat{\delta}_{cr}$ denote the sample analogue of δ_{cr}. We then estimate:

$$\hat{T}_{cr} = \rho_r + \mu_r \hat{\delta}_{cr} + \varepsilon_{cr}. \tag{3}$$

where ρ_r are indicator variables for years relative to the move year. We include these to allow for the possibility of differential trends for movers relative to non-movers that are independent of the mover's origin or destination. We weight each observation (c,r) by the total number of movers $|\mathscr{I}_c|$ in cohort c. We compute confidence intervals via 50 iterations of the Bayesian bootstrap procedure. Finally, we note that measurement error in our risky use outcome—a binary variable—would result in attenuation bias (Hausman, Abrevaya, and Scott-Morton 1998).

4.2 Interpretation

We show that under the following assumptions, the event study coefficients $\hat{\mu}_r$ in equation (3) have a precise interpretation as the weighted average of the scaled treatment effects T_{cr}/δ_{cr}. For simplicity, we focus in this section on the case where $\rho_r = 0$.

Assumption 1. Conditional Parallel Trends

$$\mathbb{E}_{i \in \mathscr{I}_c^X}\left[y_{ir}\left(\mathbf{h}_{cr}^0\right) - y_{i,-1}\left(\mathbf{h}_{cr}^0\right)\right] = \mathbb{E}\left[\hat{m}_{cr}^X - \hat{m}_{c,-1}^X\right]$$

for all c, X, and $r > -1$.

Assumption 1 is a standard (implicit) assumption in much of the prior work conducting panel mover analyses (e.g., Finkelstein, Gentzkow, and Williams 2016). It adapts the conditional parallel trends assump-

[18] We deviate from Finkelstein, Gentzkow, and Williams (2016) in allowing these origin-destination differences to be period-specific. This reflects the importance of dynamics in our context: differences in risky use between pairs of areas evolve over our sample period, as the opioid epidemic was not in steady state.

[19] This term is formally defined by $\bar{y}_{d(c)r} - \bar{y}_{o(c)r} - \mathbb{E}_{i \in \cup_c \mathscr{I}_c}\left[\bar{y}_{d(c(i))r} - \bar{y}_{o(c(i))r}\right]$. This demeaning eases the formal interpretation of the event study coefficients. It makes little difference in practice since, as we will see in Figure 3 below, the mean of $\hat{\delta}_{cr}$ in our sample is very close to zero.

tion often found in differences-in-differences settings (e.g., Heckman, Ichimura, and Todd 1997; Abadie 2005) to an event study setting. Assumption 1 will hold under our model if the distribution of both η_i and pre-move addiction status is as good as random conditional on our covariates X. In our structural analysis, we relax Assumption 1 and allow explicitly for both selection in the distributions of η_i and for $\rho_r \neq 0$.

Two other assumptions are implicit in our formulation. First, we rule out anticipatory effects by assuming that potential outcomes $y_{ir}(\cdot)$ only depend on current and past locations. Second, we also make a stable unit treatment value assumption by allowing $y_{ir}(\cdot)$ to depend only on the location history of individual i.

Proposition 1. *Under Assumption 1, $\hat{T}_{cr}$ is an unbiased and consistent estimator of T_{cr}. Furthermore, the event study coefficients $\hat{\mu}_r$ from equation* (3) *are consistent estimators of*

$$\mu_r = \sum_c w_c \frac{T_{cr}}{\delta_{cr}}, \tag{4}$$

where $w_c \equiv \frac{|\mathscr{I}_c| \cdot \delta_{cr}^2}{\sum_{c \in \mathscr{C}_m} |\mathscr{I}_c| \cdot \delta_{cr}^2}$ is the weight given to cohort c and $\sum_c w_c = 1$.

Proof. See Appendix D.

The event study coefficients μ_r thus capture a weighted average of cohort treatment effects T_{cr} relative to the average difference in outcomes between their origin and destination. Scaling by the denominator δ_{cr} gives the event study coefficients an interpretation as the average share of the gap in risky use between origin and destination that is closed as a result of the treatment effect. The weights w_c are all between 0 and 1; they are increasing in the number of individuals in a cohort and the difference in rates of risky use between their origin and destination. Weighting by the difference in risky use rates addresses extreme heteroskedasticity that would otherwise arise due to some cohorts having small values of the denominator δ_{cr}. If relative treatment effects were homogeneous across all individuals (i.e., $\frac{T_{cr}}{\delta_{cr}}$ was equal to a constant R_r) and errors in risky use were $i.i.d.$, $\hat{\mu}_r$ would be the efficient estimator of R_r.[20] □

To build intuition for what we can learn from the pattern of event study coefficients, we consider a population consisting of a single cohort c with origin location $o(c)$ and destination $d(c)$, and we assume that movers are as good as randomly selected from among individuals in their origin, conditional on X. Recall from equation (2) that T_{cr} must be zero in years prior to the move ($r < 0$), and thus μ_r must be

[20] To see this, suppose $T_{cr}/\delta_{cr} = R_r$ for all c. We can then write

$$\hat{T}_{cr} = \rho_r + R_r \delta_{cr} + \varepsilon_{cr},$$

where the errors $\varepsilon_{cr} = \hat{T}_{cr} - \mathbb{E}_{i \in \mathscr{I}_c}\left[\hat{T}_{cr}\right]$ are conditionally mean zero by Proposition 1. If the underlying individual-level errors in risky use are $i.i.d.$, the covariance matrix of ε_{cr} will have zero off-diagonal elements and diagonal elements equal to $1/|\mathscr{I}_c|$. Thus, the weighted OLS regression in equation 3 is the efficient estimator of the coefficient R_r.

equal to zero as well in those years. In years following the move we have $T_{cr} = \overline{a}_{cr}\gamma_{d(c)r} - \overline{a}_{o(c)r}\gamma_{o(c)r}$ and $\delta_{cr} = \overline{a}_{d(c)r}\gamma_{d(c)r} - \overline{a}_{o(c)r}\gamma_{o(c)r}$.[21] Thus, we can rewrite μ_r for $r > 0$ as

$$\mu_r = \frac{\overline{a}_{cr}\gamma_{d(c)r} - \overline{a}_{o(c)r}\gamma_{o(c)r}}{\overline{a}_{d(c)r}\gamma_{d(c)r} - \overline{a}_{o(c)r}\gamma_{o(c)r}}. \tag{5}$$

We consider four special cases.[22]

Case 1: Only Person Effects Suppose, first, that neither availability effects γ_{jr} nor transition probabilities π_{jr} vary between the origin and destination. Then all geographic variation is due to person effects. Rates of risky use differ across places only because of differences in the shares of addiction ($\overline{a}_{o(c)r}$ and $\overline{a}_{d(c)r}$) and the distribution of addiction propensities η_i. In this case, we have $\gamma_{d(c)r} = \gamma_{o(c)r}$ and also $\overline{a}_{cr} = \overline{a}_{o(c)r}$ (since addiction rates do not depend on location). Thus, $\mu_r = 0$ for all $r > 0$, and the event study plot will be flat at zero with no jump on move. Panel A of Figure 2 illustrates this case.

Case 2: Only Availability Place Effects Next, consider the case where availability effects γ_{jr} differ between the origin and destination but the transition probabilities π_{jr}, the shares of addiction ($\overline{a}_{o(c)r}$ and $\overline{a}_{d(c)r}$), and η_i do not. Now rates of risky use differ across places only because of differences in the ease with which individuals who are addicted can engage in risky use, and all geographic variation is due to variation in place effects on availability. In this case, we have $\overline{a}_{cr} = \overline{a}_{o(c)r} = \overline{a}_{d(c)r}$, since the evolution of addiction rates does not depend on location, and thus, $\mu_r = 1$ for all $r > 0$. The event study plot will jump from zero to one on move and remain flat at one thereafter. Panel B of Figure 2 illustrates this case.

Case 3: Person and Availability Effects Next, we combine the above two cases and allow γ_{jr}, the shares of addiction $\overline{a}_{o(c)r}$ and $\overline{a}_{d(c)r}$, and η_i vary between the origin and destination, but continue to assume the transition probabilities π_{jr} do not vary across locations. Then we have $\overline{a}_{cr} = \overline{a}_{o(c)r}$ but $\overline{a}_{o(c)r} \neq \overline{a}_{d(c)r}$. In this

[21] As discussed in Section 2, a mover is in their origin in relative year -1 and in their destination in relative year 1. In relative year 0, however, the individual may be either in their origin or destination. In discussing how the model can be used to interpret these event study coefficients, we abstract away from this empirical reality and instead assume that movers are in their origin during relative year 0 and instantaneously move before relative year 1. In the visualization of our various examples in Figure 2, we replicate uncertainty around the mover's location in relative year 0.

[22] These cases highlight the way that our model extends previous work separating place-based and person-based decompositions. Addiction creates two additional degrees of dynamics in our model that capture important aspects of the opioid crisis. First, relative to previous work—such as Finkelstein, Gentzkow, and Williams (2016)—our place-based parameters (γ_{jt}, π_{jt}) are allowed to change over time. This is reminiscent of the drift in models such as Chetty et al. (2014). However, rather than being motivated by a notion of auto-correlation, our flexibility is motivated by the dynamics of the setting. Second, the addiction channel naturally generates dynamics in treatment effects. We view the tight connection between a model of state dependence with binary states and an event-study empirical design as an additional contribution of the paper.

case, we can write the event-study coefficient for $r > 0$ as

$$\mu_r = \frac{\overline{a}_{o(c)r}\left(\gamma_{d(c)r} - \gamma_{o(c)r}\right)}{\left(\overline{a}_{d(c)r} - \overline{a}_{o(c)r}\right)\gamma_{d(c)r} + \overline{a}_{o(c)r}\left(\gamma_{d(c)r} - \gamma_{o(c)r}\right)}.$$

Here, the first term in the denominator is due only to differences in person-specific (addiction) factors between origin and destination, and the second term is due only to differences in place-specific availability. The event-study coefficient gives the share of the overall difference in outcomes due to the latter. Provided that the sign of the availability effect $\left(\gamma_{d(c)r} - \gamma_{o(c)r}\right)$ is the same as the sign of the difference in addiction rates $\left(\overline{a}_{d(c)r} - \overline{a}_{o(c)r}\right)$, this share will lie between zero and one. The event study plot will show an on-impact jump equal to this share and, provided that the ratio of addiction shares $\frac{\overline{a}_{d(c)r}}{\overline{a}_{o(c)r}}$ and availability effects $\frac{\gamma_{d(c)r}}{\gamma_{o(c)r}}$ are approximately constant, have little or no trend following the move. Panel C of Figure 2 illustrates this case.

Case 4: Adding Place-Specific Transitions to Addiction Finally, we generalize the previous case to now also allow transition probabilities π_{jr} to also vary between the origin and destination. This means that $\overline{a}_{cr}$ will differ from both $\overline{a}_{o(c)r}$ and $\overline{a}_{d(c)r}$, and we cannot simplify the expression for μ_r beyond equation (5).

However, we can build some intuition by thinking about how $\overline{a}_{cr}$ will evolve. Consider, for simplicity, the case where the destination has higher risky use rates due to both availability and addiction: $\gamma_{d(c)r} > \gamma_{o(c)r}$, $\pi^{+}_{d(c)r} > \pi^{+}_{o(c)r}, \pi^{-}_{d(c)r} < \pi^{-}_{o(c)r}$, and $\overline{a}_{d(c)r} > \overline{a}_{o(c)r}$.

Note, first, that immediately after the move $\overline{a}_{cr}$ will be close to $\overline{a}_{o(c)r}$. This means that the on-impact jump in the event-study plot will have a similar interpretation to Case 3, with the jump roughly equal to the share of differences in outcomes that are due to place-specific availability factors. Next, consider the changes in $\overline{a}_{cr}$ post-move. If most variation in addiction rates is due to person-specific factors that do not change on move, then $\overline{a}_{cr}$ will remain close to $\overline{a}_{o(c)r}$ and the event study plot will remain flat post-move as in Case 3. If a large share of variation in addiction rates is due to the place-specific terms π_{jr}, $\overline{a}_{cr}$ will gradually increase following the move from a level close to $\overline{a}_{o(c)r}$ in the direction of $\overline{a}_{d(c)r}$. This positive post-move trend in the event study plot is a signature of variation in π_{jr} playing an important role. If that variation is sufficiently important that the addiction rates among movers eventually converge to the destination level $\overline{a}_{d(c)r}$, the event study coefficients μ_r will eventually converge toward 1. Panel D of Figure 2 illustrates this case.

4.3 Results

We begin by examining the distribution of the origin-destination gaps in risky use rates across individual movers and years. This is shown in Figure 3. The mean value is close to zero and the distribution is roughly symmetric, implying that moves from states with low-rates of risky use to states with high rates of risky use are as common as moves from states with high rates of risky use to states with low rates. The standard deviation is roughly two percentage points, which is large relative to the overall mean rate of risky use among movers of approximately four percentage points shown in Table 1.

Figure 4 shows our aggregate event-study results. The figure plots the coefficients $\hat{\mu}_r$ from equation (3). The plot shows little systematic trend pre-move, which is supportive of our identifying assumption that matched non-movers serve as valid controls. It also shows two distinct features that we might expect based on Case 4 in Section 4.2 above: an immediate jump in risky use upon move and gradual post-move convergence. As discussed, the magnitude of the jump—approximately 0.28—provides a rough measure of the share of the difference between a typical origin and destination attributable to differences in opioid availability to addicted individuals. This might reflect place-specific factors such as the presence of "pill mill" pain clinics or prescription monitoring programs intended to limit risky use. In addition, the significant post-trend that we see may suggest a significant role for place-specific factors that affect transitions in and out of addiction. This might reflect place-specific factors such as the propensity of local physicians to initiate opioid prescriptions. Together, these features suggest that both opioid availability and addiction transitions play a significant role in driving cross-sectional variation in risky use.

Column 1 of Table 2 provides a quantification of these estimates, summarizing the average values of $\hat{\mu}_r$ at various time horizons r relative to the move. One year after moving, we estimate that this coefficient is 0.275 (standard error = 0.033). Five years after moving, this estimate grows to 0.573 (standard error = 0.053). To illustrate these dynamics, consider locations with a stable 3.5 percentage point difference between non-movers in the rate of risky use—approximately equivalent to the difference between the 20th percentile state-year and the 80th percentile state-year in average rates of risky use over our full sample period. A move between these two locations would be associated with an immediate increase in the likelihood of risky use of approximately 1 percentage point in the year after moving, or about 30 percent of the cross-sectional gap between these areas. On average, the likelihood of risky use would continue increasing by about 0.30 percentage points each subsequent year after the move, with approximately 60 percent of the cross-sectional gap closed by five years after the move.

Figure 5 provides another way of visualizing how the propensity for risky use changes around moves. It

shows a binscatter of the average change in risky use rates over the one to five years post-move compared to the one to five years prior to move against the origin-destination difference $\hat{\delta}_{ct}$ in risky use rates. The results show a clear relationship between the size of the move and changes in the prevalence of risky use, with a significant and positive slope (0.27). This plot also suggests that the relationship between the size of the move and the change is roughly linear and symmetric for moves between places with higher and lower risky use rates, consistent with an assumption of roughly constant scaled treatment effects $\frac{T_{cr}}{\delta_{cr}}$. We also find roughly symmetric effects when we estimate equation (3) for moves to areas with higher and lower risky use rates separately, with both types of moves showing a change in the individuals' probability of risky use immediately when they move (Appendix Figure A.6).

Finally, to get a sense of how movers are selected in their average risky use rates, we examine the event study estimates when we no longer adjust for any fixed individual differences in propensity for risky use. Specifically, instead of defining $\hat{T}^X_{cr}$ to be the average of the difference in differences $(y_{ir} - y_{i,-1}) - \left(\hat{m}^X_{cr} - \hat{m}^X_{c,-1}\right)$ among the set $\mathscr{I}^X_c$ of movers in cohort c with characteristics X, we define the dependent variable as $\left(y_{ir} - \hat{m}^X_{cr}\right)$. Appendix Figure A.7 shows the results of this alternative "levels" specification. We continue to find little systematic pre-move trend, and a similar post-move pattern of coefficients. However, the level of the pre-move coefficients is now about 0.10, indicating that there is potential selection in where movers move. In particular, the results indicate that movers to destinations with a higher propensity for risky use than in their origin tend, prior to their move, to have a higher propensity for risky use than matched non-movers in their origin. We parameterize our model below to explicitly allow for such selection within the context of the model.

Heterogeneity by prior opioid use The model described in Section 3 predicts differential treatment effects by the addiction status of movers. Changes in availability, for example, would directly affect addicted individuals immediately upon move, but would have no immediate effects on those who are not addicted. Differences in place-based transitions to and from addiction would affect all individuals, and so lead to post-trends in the event study for both those who are addicted and those who are not. The aggregate event study in Figure 4 captures an average of these effects.

While we cannot directly observe addiction states in a given year, we explore these predictions empirically by proxying for addiction status with an individual's prior history of opioid use. This dovetails with medical literature studying how addiction varies with the opioid histories of patients (Paulozzi et al. 2012; Edlund et al. 2014), and exploring the impacts of drug supply among patients without any previous opioid use (Shah, Hayes, and Martin 2017; Brat et al. 2018; Jeffery et al. 2018). Furthermore, the heterogeneous

effects of moving on each group highlight the importance of considering both the relative effectiveness of policies in settings with larger and smaller existing stocks of addiction as well as their potentially differential effects on prior users and opioid-naives.

Figure 6 shows the results from estimating equation (3) separately for movers with at least some opioid utilization in the year prior to move ("prior users") and movers without opioid use in the year prior to move ("opioid naives"). These results are summarized in Table 2 (columns 2 and 3) and are consistent with the predictions of the model. Prior users show a larger immediate jump upon move, while there is little or no discrete change upon move for opioid naives. In the subsequent post-move years, relative risky use rates for both prior users and opioid naives increase gradually, consistent with high risky use areas causing more transitions to addiction and fewer transitions from addiction.

Robustness In Appendix C, we show that these results on the changes in risky prescription opioid use around moves—both in aggregate and separately based on prior opioid use—are robust to several alternative specifications. These include using county or commuting zone as the unit of analysis instead of state, using alternative measures of risky prescription opioid use, omitting relative year fixed effects (ρ_r), and removing our conditioning on covariates (X). They also continue to hold when we estimate the event study on a balanced panel of individuals whom we see with Part D coverage for at least three years before and after move, if we restrict to individuals who had Part D coverage for the full year for all years they appear in the sample, or if we restrict to individuals who are alive throughout our sample period. They are robust to removing the three most common destinations in our sample (Florida, Texas, and California) as well as excluding patients in hospice or being treated for cancer. Finally, we find similar results when we re-estimate our main specifications in the full sample of elderly Medicare recipients.

5 Model Parameterization and Estimation

The evidence in Section 4 suggests an important role for place-specific factors affecting both opioid availability and addiction transitions. To make quantitative statements about the relative importance of various channels in driving risky opioid use, and to assess the role of specific policies, we now turn to parameterizing and estimating the dynamic model of opioid addiction and risky opioid use developed in Section 3. We use the same sample as in Section 4, except that we omit observations from the move year $(r = 0)$, since the enrollee may spend time in both their origin and their destination in that year (see Appendix Figure A.2).

5.1 Parameterization

To fix the initial conditions of the model, we define $t = 0$ to be 1995, the year before the introduction and marketing of OxyContin, an event that is "often seen as the trigger event for the current opioid crisis" (Alpert et al. 2022; Maclean et al. 2022). Our assumption that the initial addiction state is $a_{i0} = 0 \, \forall \, i$ then implies that there was no addiction prior to 1996, and that the addiction states evolved according to our model thereafter.

For computational feasibility, we impose some additional parametric assumptions. We assume that the baseline addiction parameters are time constant in the absence of policy effects so that $\pi^{+}_{jt0} = \pi^{+}_{j0}$ and $\pi^{-}_{jt0} = \pi^{-}_{j0}$. We further assume that the baseline availability parameter γ_{jt0} can be decomposed into time and location fixed effects, $\gamma_{jt0} = \gamma_j + \tau_t$. This ensures that geographic differences in opioid availability remain constant (in proportions) throughout our sample while allowing for shifting national attitudes and guidelines to affect availability across all states. Finally, we abstract from any within-location heterogeneity in the addiction transition parameters η_i among non-movers, so that $\eta_i = \eta_j$ for all i such that $i \in \mathscr{J}_j$.

5.2 Policy effects

Motivated by the evidence discussed in Section 2.3, we focus our policy analysis on pill mill laws. We allow for the implementation of pill mill laws to affect both the propensity of individuals in an area to transition into addiction (via π^{+}_{jt}) and the availability of opioids to those who are addicted (via γ_{jt}).[23] This flexibility is consistent with recent evidence on the effects of pill mill policies. Although these regulations were intended to regulate inappropriate prescribing patterns and to target institutions associated with easily-accessible, frequent, and large prescriptions—which we might think of as closely tied to the availability channel—the literature on pill mill policies suggests that these policies likely had effects on both risky use and potentially appropriate use that could lead to transitions into addiction (Rutkow et al. 2015; Lyapustina et al. 2016; Kaestner and Ziedan 2023).

Based on the evidence discussed in Section 2.3, we allow for these policies to be gradually implemented by assuming that

$$\begin{aligned} \pi^{+}_{jt} &= \left(1 - \zeta^{\pi}_{jt}\right) \pi^{+}_{jt0} \\ \gamma_{jt} &= \left(1 - \zeta^{\gamma}_{jt}\right) \gamma_{jt0}, \end{aligned}$$

where π^{+}_{jt0} and γ_{jt0} are the baseline addiction and availability parameters that individuals would face in the absence of a pill mill policy and $\zeta^{\pi}_{jt}, \zeta^{\gamma}_{jt} \in [0,1]$ are the extent to which these parameters are changed due

[23] For tractability, we do not allow them to change the probability π^{-}_{jt} of transitioning out of addiction.

to the policy as of year t in state j. We assume that policy impacts are weakly increasing over their first five post-enactment years, by which point they are fully phased in. More precisely, letting t_j^* denote the year the policy is enacted in state j, we assume that: (i) $\zeta_{jt}^{\pi}, \zeta_{jt}^{\gamma} = 0$ for $t < t_j^*$ and for states that never enact policies, (ii) $\zeta_{jt}^{\pi}, \zeta_{jt}^{\gamma}$ attain their maximal values $\zeta_{max}^{\pi}, \zeta_{max}^{\gamma}$ for $t \geq t_j^* + 5$, and (iii) ζ_{jt}^{π} and ζ_{jt}^{γ} are weakly increasing in t for $t \in \left[t_j^*, t_j^* + 4\right]$, with $\zeta_{jt}^{\pi} = share_{t-t_j^*}\zeta_{max}^{\pi}$ and $\zeta_{jt}^{\gamma} = share_{t-t_j^*}\zeta_{max}^{\gamma}$ for $share_{t-t_j^*} \in [0,1]$ a set of monotonically increasing weights that capture the share of the full policy that is in force $t - t_j^*$ years after policy enactment. In Section 6.2 we show that our counterfactuals are robust to alternative timing assumptions for gradual policy implementation.

5.3 Allowing for relative year effects and selection effects

We extend our baseline model in Section 3 to directly allow for relative year fixed effects and potential selection among movers.

We allow for the risky use rates of addicted movers to differ from non-movers through proportional shifts in risky use rates that evolve arbitrarily in years around the move; we denote these relative year fixed effects by ρ_r. These shifts capture changes in use associated with moving *per se*, such as disruption of supply networks, changes associated with family stress, and so on, as well as correlation between the timing of moves and positive or negative shocks to opioid demand (provided these do not vary systematically depending on the origin or destination).

Combining these relative year fixed effects with our previous assumption that the baseline availability parameters are additively separable implies that risky use among movers who are addicted during a given year is given by

$$\gamma_{cr} = \begin{cases} \gamma_{o(c)} + \tau_{t(c,r)} + \rho_r & \text{if } r < 0 \\ \gamma_{d(c)} + \tau_{t(c,r)} + \rho_r & \text{if } r > 0. \end{cases}$$

Next, we address potential selection explicitly by allowing for the distribution of person-specific addiction transitions to vary systematically with both the mover's origin and their destination. Specifically, we assume that the person-specific parameters for each move cohort c are drawn from some convex combination of the distributions of person-specific factors in their origin state $o(c)$ and destination state $d(c)$. The distribution of person-specific parameters in each cohort is therefore given by

$$\eta_c \sim (1-s) \cdot \eta_{o(c)} + s \cdot \eta_{d(c)}.$$

The selection parameter s governs the extent to which movers are selected to resemble individuals in their destinations. Thus, the predicted shares of addiction in our model for each cohort c evolve according to

$$\bar{a}_{cr} = \begin{cases} (1-\pi^{-}_{o(c)}-\eta^{-}_{c})\bar{a}_{c,r-1} + (\pi^{+}_{o(c),t(c,r)}+\eta^{+}_{c})\bar{a}_{c,r-1} & \text{if } r<0 \\ (1-\pi^{-}_{d(c)}-\eta^{-}_{c})\bar{a}_{c,r-1} + (\pi^{+}_{d(c),t(c,r)}+\eta^{+}_{c})\bar{a}_{c,r-1} & \text{if } r>0 \end{cases}$$

where $\eta_c = (1-s)\cdot\eta_{o(c)} + s\cdot\eta_{d(c)}$. We do not use risky use in the year of the move as a moment in our estimation, and assume in our baseline specification that no addiction transitions occur during this time. We show in Section 6.2 that allowing for addiction transitions to occur during the year of the move does not substantially affect our results, and that assuming no selection ($s=0$) also does not affect our results.

5.4 Estimation

Our sample includes 16,612 mover cohorts—defined by unique combinations of origin state, destination state, and move year—and 714 non-mover state-years. This results in 125,921 mover moments ($\hat{y}_{cr}$) and 714 non-mover moments ($\hat{y}_{jt}$), which are the sample analogues of the cohort risky use rates $\bar{y}_{cr}$ and non-mover risky use rates $\bar{y}_{jt}$.[24]

We estimate the parameters of our model by Generalized Method of Moments (GMM), with moment conditions defined for movers at the level of the cohort by relative year and non-movers at the state-year level. We weight each moment by the number of beneficiary-years in the sample used to construct the moment. However, because non-movers make up the vast majority of our sample, despite movers providing most of the critical variation for our model, we re-weight non-mover moments in aggregate so that movers are up-weighted by a factor of approximately five.[25] We also restrict our parameter space so that all transition probabilities are bounded between zero and one and so that our combination of parameters produces a weighted global mean share of addiction across all years of 10 percent.[26] Appendix E presents more de-

[24]The evidence from Section 4 suggests that including demographic covariates does not affect the basic event study patterns. For computational ease, we therefore do not further separate our sample into additional moments by demographic. Appendix Table A.3 presents additional details about the moments used for estimation.

[25]While there are approximately 28 times more non-movers in our sample, we re-weight movers in aggregate so that this aggregate weighting ratio is $\sqrt{\#\text{Non-mover Years}}$ to $\sqrt{\#\text{Mover Years}}$.

[26]The share of the population who is addicted to opioids is not directly observed in the data. We chose 10 percent as a reasonable approximation; it is between the share of the population in the 2011 National Survey on Drug Use and Health that reports non-medical prescription pain reliever use over a 12-month period (4 percent) and the share that reports such use 'ever' (13 percent) per authors' calculations. We show in Appendix G that other values do not substantially affect our results; indeed, the global mean share of addiction essentially serves as a normalization because the relative overall averages of addiction and availability parameters are not separately pinned down in our model. For example, doubling the shares of individuals who are addicted and halving all

tails on our estimation and presents Monte Carlo results suggesting that the finite-sample properties of our estimator and key counterfactuals are reasonable.

5.5 Identification

We offer intuition for the key features of the data that provide identifying variation for our estimates. This discussion builds on the intuition developed in Section 4.2 above.

Person-specific factors are constant in time and do not change upon move. Therefore, the distribution of η_i will be informed by patterns in risky use rates that are common to movers from a given origin regardless of their destination. Our estimate of the selection parameter s is related to the way pre-move patterns of risky use differ across movers with different destinations.

In contrast, place-specific factors do change discretely upon move, and generate sharp changes in predicted risky use directly after moving. The estimation of differences in place-specific availability effects γ_{jt} will therefore be driven by the changes in risky use that occur immediately upon move. Similarly, place-specific addiction parameters π will be related to changes in the evolution of risky use rates in the years post-move.

A few other parameters are worth commenting on. To see how we identify relative year effects (ρ_r), consider two cohorts with the same origin and destination but different move years. During the same pre-move calendar year, the two cohorts will share the same origin, destination, and calendar year. Thus, they will differ only in their year relative to move, and the ratio of their risky use will drive our estimates of ρ_r.

To see how we separately identify the addiction transition parameters into (π^+ and η^+) and out of addiction (π^- and η^-) , it is useful to consider two hypothetical cohorts of movers: one cohort in which everyone in the mover population is initially addicted and another cohort in which no one is initially addicted. The cohort in which everyone is addicted would only initially be affected by transitions out of addiction, while the cohort in which no one was addicted would only initially be affected by transitions into addiction. In the more general case, transitions into and out of addiction are related to local patterns of rising and declining risky use rates that cannot be explained by national trends in availability (τ_t).

Finally, our non-mover moments contribute to the identification of the pill mill policy effects (ζ^{π}_{jt} and ζ^{γ}_{jt}). Although pill mill policies also generate changes in predicted risky use for movers, the patterns of non-movers before and after the policy implementation provide additional power for estimating the direction and magnitude of the policy effects. Our assumption that the policy is fully implemented within five years after enactment also allows us to separate the sources of dynamic treatment effects. Dynamics implied by gradual

availability parameters would leave predicted risky use at a point in time unchanged.

implementation occur only for the first five years of the policy, while longer-run changes in dynamics of risky use are driven only by the policy's effects on addiction.

6 Model Estimates

6.1 Model fit and parameter estimates

Model fit The model fits the data well. In particular, the simulated moments from our estimated model closely match empirical moments in the data (see Appendix Figures A.8, A.9, and A.10). In addition, the estimated model can reproduce our reduced form findings, which were not explicitly targeted in estimation. Figure 7 shows that we approximately match the main mover event study from Figure 4. Appendix Figure A.11 shows that we also match the estimated impact of pill mill regulations based on their staggered introduction across states (originally presented in Appendix Figure A.3, Panel B).[27]

Parameter estimates The panels of Figure 8 show the distribution across states of factors affecting availability, place-based addiction transitions, and person-based addiction transitions (among non-movers), as well as the correlation of these factors with predicted risky use rates.[28] They suggest that the three sets of parameters differ substantially in their geographic distributions. For example, availability (γ_j) plays a large role in driving risky use in North Dakota and Vermont relative to addiction, and it plays a larger relative role in the West and the Northeast. In contrast, place-specific and person-specific addiction factors (π_j and η_i) play relatively larger roles in driving risky use in the South and Appalachia. Figure 8 also shows that our estimated parameters are all highly correlated with overall risky use rates.

Table 3 summarizes the parameter estimates corresponding to the impact of the introduction of pill mill policies. The estimates suggest that the largest effects of pill mill policies were on addiction channels. In particular, row (1) suggests that once fully implemented, pill mill laws cut the probability of transitioning into addiction by nearly 34% in implementing states. On the other hand, row (2) suggests that the impact of pill mill policies on availability were small in comparison, although we cannot rule out moderate effects on the order of 5 - 10% in implementing states. The sizable effect on non-addicted individuals is consistent

[27] We include simulated standard errors in both Figure 7 and Appendix Figure A.11 to provide a sense of the uncertainty of the estimators given our actual sample sizes. We also note that, as discussed in Section 5.5, variation from movers and non-movers identify distinct parts of our model. Consistent with this, we find that increasing the weight on mover moments in our model improves the fit of our simulated mover event study while increasing the weight on non-mover moments improves the fit of our estimated pill mill regulation effects.

[28] In particular, we examine geographic variation in pre-policy parameters (γ_j, π_j, and η_i). We omit the time subscript as our parameterizations in Section 5.1 assume that in the absence of policy affects, baseline addiction parameters are time constant in the absence of policy effects and geographic variation and temporal variation in availability are separable.

with evidence from other opioid-related policies. For example, Alpert, Dykstra, and Jacobson (2024b) find that that Kentucky's landmark PDMP introduced physician hassle costs that meaningfully reduced opioid prescribing to both prior users and opioid-naive patients. Consistent with the evidence of implementation delays discussed in Section 2.3, rows (3) - (7) suggest the pill mill policy impacts begin in the year after implementation and grow considerably over time.

Area-level correlates of parameter estimates Although we only explicitly model one observable place-specific characteristic (the implementation of state-specific pill mill policies), we also explore how our parameter estimates correlate with other observable area-level characteristics. This provides some suggestive evidence on other factors behind our addiction and availability parameters, although it does not have a direct causal interpretation.

In order to exploit finer geographic variation in these characteristics, for these analyses we hold our estimated policy parameters fixed and re-estimate our other model parameters at the commuting-zone level, focusing on the largest commuting zones . This allows us to look at the correlations across 91 geographic areas, rather than across 50 states.[29] Figure 9 presents bivariate area-level correlations between various parameter estimates (shown across the different vertical panels) and area-level characteristics (shown across the rows); for reference, the first panel shows the correlation of the area-level characteristics with the average rate of risky use for our sample in that area. The first row considers the average prescription opioid dose given after surgery. Intuitively, giving high doses of opioids to surgery patients is a leading pathway for transitions into addiction. Consistent with this, the results indicate that places where surgery patients receive higher average doses have higher rates of risky use, and our estimates suggest that this reflects their higher rates of place-based transitions into addiction, but not different availability or person-based addiction factors. This finding is consistent with Eichmeyer and Zhang (2022), who leverage local variation in physician opioid prescribing among veterans with emergency department visits and find strong effects on downstream transitions to opioid dependence. The second row considers state-level triplicate prescribing restrictions.[30] Places with these restrictions have lower risky use rates, and our estimates suggest this is again driven by a strong relationship with place-based addiction factors. This pattern is consistent with Alpert et al. (2022),

[29] Appendix C presents the results from estimating the model at the commuting-zone level and shows that the model fit is good and that the relative importance of various parameters is similar to our baseline, state-level model. Appendix Appendix H presents detailed definitions, data sources, and summary statistics for the various correlates we examine. We also conduct the same exercise at the state-level in Appendix Figure A.12, where we find broadly similar, albeit noisier, results. For the commuting-zone correlates, we present a wider set of demographic correlates in Appendix Table A.4.

[30] As Alpert et al. (2022) discuss, these policies were in place long before the beginning of the opioid epidemic in the 1990s. For the purpose of this cross-sectional analysis, we do not examine correlations with policies that were passed during our sample period such as PDMPs or prescribing limits, as such correlations could reflect both policy impacts and endogeneity of policy adoption.

who show that triplicate prescribing restrictions were associated with less OxyContin-related marketing.

The middle rows of the figure indicate that places with poorer quality of healthcare, as measured by a higher rate of hospitalizations for ambulatory care sensitive conditions (ACSC)—i.e., hospital admissions that could potentially have been avoided through better quality outpatient care—have lower rates of risky use, mainly because of lower estimated availability. This may reflect more constrained health care supply in these lower-quality areas. Finally, the bottom panel shows that places with lower manufacturing employment, higher college-education shares, and higher incomes tend to have higher availability parameters (consistent with greater overall health care supply) but lower place-based addiction parameters (consistent with the general negative relationship between socio-economic status and addiction rates); these opposing effects largely cancel out, leading our economic characteristics to be mostly uncorrelated with risky prescription use in aggregate.

6.2 Drivers of Risky Use

Table 4 presents evidence on the quantitative importance of various model parameters for risky use. It shows the average impact on risky use in 2019 of a one standard deviation reduction in different channels (holding all others constant) over our 2006-2019 study period. Consistent with our results from Section 4, we find quantitatively large roles for geographic variation in both person- and place-based channels in driving risky use. A one standard deviation reduction in the person-specific addiction propensities (η_i) would have reduced risky use by approximately 33 percent over our fourteen-year period (row 1). A one standard deviation reduction in both place-specific addiction (π_j) and availability (γ_j) would have had a similar effect (row 2).

The results also indicate a substantially larger role for addiction parameters than availability parameters in driving risky use over our study period. The effect of reducing place-specific addiction parameters by one standard deviation is more than twice as large as the effect of reducing place-specific availability (rows 2a and 2b). Reducing both sets of addiction parameters (π_j and η_i) by one standard deviation would reduce risky use by more than 50 percent (row 3), while reducing availability alone would only reduce it by 14 percent (row 2b).

However, these 'point in time' results mask important heterogeneity in the time path of the two distinct place-based channels. To illustrate this, Figure 10 shows the time path of the effects of a one standard deviation reduction in place-based addiction transitions and a one standard deviation reduction in availability. In the first year, we estimate that average risky use rates would have been 20 percent lower if availability were targeted, while targeting place-based addiction has no immediate impact on risky use. With each passing

year, however, the relative impact of targeting place-based addiction transitions grows, and the relationship is ultimately reversed. In the third year after the hypothetical policies are implemented, the reductions in risky use are nearly equal across policies. By the end of our study period, the impact on risky use of the reduction in place-based addiction transitions is nearly twice as large as the reduction in availability.

Alternative Specifications Appendix Table A.5 shows that these quantitative conclusions about our parameter estimates are generally similar across a variety of alternative samples as well as extensions or modifications to our model framework. In particular, the larger role for addiction relative to availability in driving risky use is robust across the alternative specifications. Our finding that both person-based and place-based parameters are quantitatively important in driving risky use is also robust across alternative specifications, although the relative magnitude of person- and place-based factors varies across the specifications.

The first row replicates the baseline results from Table 4. In rows (2) and (3), we consider two alternative samples. First, we re-estimate our model by using a subset of large commuting zones and re-defining movers and non-movers using these alternative geographic boundaries. Second, we replicate key analyses on an alternative sample of elderly Medicare (65+) enrollees. As in our baseline sample, we find that geographic variation in addiction plays a larger role than variation in availability. Compared to our baseline specification, in these alternative specifications person-based parameters appear to play a slightly larger role relative to place-based parameters.

The rest of the table describes results from alternative modeling specifications, which we describe in more detail in Appendix G. In rows (4) - (6) we consider a set of models which allow place-based addiction factors and person-based addiction factors to interact non-linearly through a CES-style specification and an aggregator α; specifically, we allow an interaction that is, respectively, a multiplicative function, a minimum function, and a maximum function.[31] The results from the multiplicative function are quite similar to our baseline results, while the more extreme specifications—the minimum and the maximum function—vary more, but consistently indicate large roles for the addiction parameters relative to availability in driving risky use.

In row (7), we estimate a model specification in which we allow for individuals who are not addicted to

[31] In particular, this corresponds to the following (more general) addiction transition equation.

$$Pr\left(a_{i,t}=1|j\left(i,t\right)=j\right)=\begin{cases}1-\left(\left[\pi_{jt}^{-}\right]^{\alpha}+\left[\eta_{i}^{-}\right]^{\alpha}\right)^{\frac{1}{\alpha}} & \text{if } a_{i,t-1}=1\\ \left(\left[\pi_{jt}^{+}\right]^{\alpha}+\left[\eta_{i}^{+}\right]^{\alpha}\right)^{\frac{1}{\alpha}} & \text{if } a_{i,t-1}=0.\end{cases}$$

Within this CES-style specification, our baseline model corresponds to $\alpha=1$, where place-based addiction factors and person-based addiction factors are additively separable. The multiplicative, minimum, and maximum edge case specifications correspond to $\alpha\to 0$, $\alpha\to-\infty$, $\alpha\to\infty$ respectively.

engage in risky use. For tractability, we assume that this probability is constant across time and geographies, and we estimate this probability as an additional parameter in our model. This non-zero probability of risky use among non-addicted individuals means that policies aimed at reducing place-based availability and place-based addiction are mechanically likely to have smaller effects, as some share of risky use is now driven by non-addicted individuals who these counterfactuals do not affect. The relative importance of each channel, however, remains similar to our baseline model.

Finally, we also estimate specifications where: (i) we allow addiction transitions to occur during the year of the move (row 8), (ii) we shut down selection such that $s = 0$ (row 9), (iii) we consider alternative normalizations for the global share of addiction $\bar{a}$ (rows 10 and 11), and (iv) where we consider various different timing assumptions for gradual policy implementation (rows 12 and 13).

6.3 Policy Counterfactuals

Figure 11 shows the simulated impact of pill mill policies in the 11 states that adopted them. We consider the impact of the policies using the actual policy timing as well as two counterfactuals: one in which these states passed the policies at the start of our data (2006) and another in which these states passed the policies in the year before OxyContin was first introduced (1995). In each scenario, we continue to allow for implementation delays over five years, as policies become modified, strengthened, and enforced. Panel A shows the results in terms of percentage changes, while Panel B shows the results in levels.

We estimate that pill mill policies reduced the average rate of risky use between 2006 and 2019 by 4.4 percent in implementing states, with gains concentrated late in the period. If the adopting states had all implemented their pill mill regulations in 2006, these policies would have reduced risky use over this same time period by 14.9 percent. Had these policies been in place since 1995, cumulative risky use over the same period would have fallen by 30.0 percent.

Early policy implementation thus yields substantially larger reductions in risky use. We show in Appendix Figure A.13 that the bulk of these reductions are driven by the dynamic effects of targeting addiction earlier rather than the effect of avoiding implementation delays. In particular, Panel A considers the gains when we remove the implementation delay, and Panel B considers the gains when we start our model with the smaller stock of addiction generated from the 1995 implementation. Our results suggest that approximately 75 percent of the additional gains from early implementation are from targeting addiction before a stock of addiction has accumulated, while the remaining gains are from avoiding implementation delays. Our findings thus echo the analysis in Cutler and Glaeser (2021), who emphasize that the persistence of addiction creates challenges to ending the opioid epidemic once it is underway.

Potential substitution to illegal opioids An important limitation to the counterfactuals in Table 4 and Figure 11 is that it only examines risky use of *prescription* opioids. If reductions in the availability of prescription opioids cause substitution toward illegal opioids—as has been suggested by prior work (e.g. Evans, Lieber, and Power 2019; Alpert, Powell, and Pacula 2018)—the welfare implications of these counterfactuals could be attenuated or even reversed. Our data do not permit a direct examination of such potential substitution, and this is an important area for future work. But we present some suggestive evidence that any such substitution may be relatively limited.

Specifically, Table 5 shows the area-level correlation between our estimated model parameters and rates of three adverse opioid-related health outcomes that include cases due to both prescription and illegal opioids: opioid poisonings, opioid use disorders, and total drug poisonings.[32] Our estimates of addiction and availability parameters are strongly predictive of both opioid poisonings and total drug poisonings (columns 2 and 3). A one standard deviation increase in place-based addiction in a commuting zone is associated with a 0.027 percentage point increase in opioid-related poisonings each year (approximately 20 percent of the mean) and a 0.062 percentage point increase in total drug-related poisonings each year (more than 10 percent of the mean). We do not find a significant correlation with diagnoses of opioid-use disorder, which could reflect places that exert stricter control over opioid prescriptions also diagnosing opioid use disorder more aggressively.

This pattern of correlations is consistent with relatively limited substitution to illegal opioids. To see this, consider a simple model in which people may use both risky prescription opioids and illegal opioids, and the latter is at least as likely as the former to produce adverse outcomes. At one extreme, if reductions in prescription use cause one-for-one substitution to illegal use, the relationship between availability of prescription opioids and adverse outcomes should be *zero or negative*, since making prescription opioids less easily available would cause an equivalent increase in illegal use and thus weakly increase the overall likelihood of adverse outcomes. At the other extreme, if there is no substitution between prescription opioids and illegal opioids, the relationship between availability of prescription opioids and adverse outcomes should be positive; we also show that in this case, the effects of availability and addiction on adverse outcomes should be roughly proportional to their effects on risky use. We present a formal model in Appendix F that makes the logic sketched here more precise.[33] The correlations in Table 5 are more consistent with the prediction under independence than with the prediction under full substitution.

[32] As with the analysis of other area correlations in Figure 9, this analysis is done at the commuting zone level to exploit finer geographic variation.

[33] For a more developed economic model of choice between prescription and illicitly-manufactured sources of opioids for non-medical use, see Schnell (2022) and Mulligan (2024).

7 Conclusion

We explore drivers of the prescription opioid epidemic between 2006 and 2019 and the role that pill mill laws played, focusing on individuals enrolled in the Social Security Disability Insurance (SSDI) program, a population hit particularly hard by the opioid epidemic. Both our reduced form results and our estimates from a dynamic model of risky opioid use highlight important roles for both person-specific factors and place-based factors in driving risky prescription opioid use.

Our findings indicate that a one standard deviation reduction in place-specific factors would have reduced risky use by about 40 percent. They also illustrate the distinct role played by two separate place-based channels: an addiction channel which affects transitions in and out of addiction, and an availability channel which affects the ease with which addicted individuals can obtain risky prescriptions. These channels create a key temporal tradeoff: policies targeting availability are more effective in the first few years, while policies targeting place-based addiction propensities become substantially more effective after that. As Cutler and Glaeser (2021) highlight, the existing stock of addiction may require its own policies to tackle, such as designing policies to increase access to treatment facilities (Corredor-Waldron and Currie 2022).

Of course, a single policy may simultaneously impact multiple channels at once. Indeed, our estimates of the impacts of pill mill policies indicate that they decrease both availability and transitions into addiction. We estimate that, relative to a world without these policies, by 2019 pill mill policies decreased risky use by about 4.4 percent in the eleven states that implemented these policies during our study period. If, counterfactually, these pill mill policies had been implemented in 1995, they would have decreased risky use by 30 percent. The substantially larger impact of earlier implementation primarily reflects the importance of the addiction channel; the effects of preventing addiction from growing compound over periods and are quantitatively large. More broadly, they illustrate how the dynamic time-path of policy impacts can depend critically on how much they impact the availability channel compared to the addiction channel.

References

Abadie, Alberto. 2005. "Semiparametric Difference-in-Differences Estimators." *The Review of Economic Studies* 72 (1): 1–19.

Aliprantis, Dionissi, Kyle Fee, and Mark E. Schweitzer. 2023. "Opioids and the labor market." *Labour Economics* 85 102446.

Alpert, Abby, Sarah Dykstra, and Mireille Jacobson. 2024a. "Hassle Costs versus Information: How Do Prescription Drug Monitoring Programs Reduce Opioid Prescribing?" *American Economic Journal: Economic Policy* 16 (1): 87–123.

Alpert, Abby, Sarah Dykstra, and Mireille Jacobson. 2024b. "Hassle Costs versus Information: How Do Prescription Drug Monitoring Programs Reduce Opioid Prescribing?" *American Economic Journal: Economic Policy* 16 (1): 87–123.

Alpert, Abby, William N Evans, Ethan M J Lieber, and David Powell. 2022. "Origins of the Opioid Crisis and its Enduring Impacts." *The Quarterly Journal of Economics* 137 (2): 1139–1179.

Alpert, Abby, David Powell, and Rosalie Liccardo Pacula. 2018. "Supply-Side Drug Policy in the Presence of Substitutes: Evidence from the Introduction of Abuse-Deterrent Opioids." *American Economic Journal: Economic Policy* 10 (4): 1–35.

Autor, David H. 2015. "The Unsustainable Rise of the Disability Rolls in the United States: Causes, Consequences, and Policy Options." In *Social Policies in an Age of Austerity: A Comparative Analysis of the US and Korea*, edited by Scholz, John Karl, Hyunpyo Moon, and Sang-Hyop Lee 107–136, Northampton MA: Edward Elgar Publishing.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103 (6): 2121–2168.

Autor, David H., and Mark G. Duggan. 2006. "The Growth in the Social Security Disability Rolls: A Fiscal Crisis Unfolding." *Journal of Economic Perspectives* 20 (3): 71–96.

Baehren, David F., Catherine A. Marco, Danna E. Droz, Sameer Sinha, E. Megan Callan, and Peter Akpunonu. 2010. "A statewide prescription monitoring program affects emergency department prescribing behaviors." *Annals of Emergency Medicine* 56 (1): 19–23.e1–3.

Barnett, Michael L., Andrew R. Olenski, and Anupam B. Jena. 2017. "Opioid-Prescribing Patterns of Emergency Physicians and Risk of Long-Term Use." *New England Journal of Medicine* 376 (7): 663–673.

Barrette, Eric, Leemore Dafny, and Karen Shen. 2023. "Do Policies to Increase Access to Treatment for Opioid Use Disorder Work?" *American Journal of Health Economics* 9 (3): 297–330.

Billings, John. 2003. "Using Administrative Data To Monitor Access, Identify Disparities, and Assess Performance of the Safety Net."Technical report, Agency for Healthcare Research and Quality.

Bohnert, Amy S. B., Marcia Valenstein, Matthew J. Bair, Dara Ganoczy, John F. McCarthy, Mark A. Ilgen, and Frederic C. Blow. 2011. "Association Between Opioid Prescribing Patterns and Opioid Overdose-related Deaths." *JAMA* 305 (13): 1315–1321.

Braden, Jennifer Brennan, Joan Russo, Ming-Yu Fan, Mark J. Edlund, Bradley C. Martin, Andrea DeVries, and Mark D. Sullivan. 2010. "Emergency Department Visits Among Recipients of Chronic Opioid Therapy." *Archives of Internal Medicine* 170 (16): 1425–1432.

Brat, Gabriel A, Denis Agniel, Andrew Beam et al. 2018. "Postsurgical Prescriptions for Opioid Naive Patients and Association with Overdose and Misuse: Retrospective Cohort Study." *BMJ* 360.

Bronnenberg, Bart J., Jean-Pierre H. Dubé, and Matthew Gentzkow. 2012. "The Evolution of Brand Preferences: Evidence from Consumer Migration." *American Economic Review* 102 (6): 2472–2508.

Buchmueller, Thomas C., and Colleen Carey. 2018. "The Effect of Prescription Drug Monitoring Programs on Opioid Utilization in Medicare." *American Economic Journal: Economic Policy* 10 (1): 77–112.

Callaway, Brantly, and Pedro H.C. Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics* 225 (2): 200–230.

Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality." *The Quarterly Journal of Economics* 128 (3): 967–1015.

Carpenter, Christopher S., Chandler B. McClellan, and Daniel I. Rees. 2017. "Economic Conditions, Illicit Drug Use, and Substance Use Disorders in the United States." *Journal of Health Economics* 52 (C): 63–73.

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): 15078–15083.

Case, Anne, and Angus Deaton. 2017. "Mortality and Morbidity in the 21st Century." *Brookings Papers on Economic Activity, Spring 2017*.

Case, Anne, and Angus Deaton. 2020. *Deaths of Despair and the Future of Capitalism*. Princeton: Princeton University Press.

CDC. 2017. "U.S. Prescribing Rate Maps." `https://www.cdc.gov/drugoverdose/maps/rxrate-maps.html`.

Cepeda, M. Soledad, Daniel Fife, Wing Chow, Gregory Mastrogiovanni, and Scott C. Henderson. 2012. "Assessing Opioid Shopping Behaviour." *Drug Safety* 35 (4): 325–334.

de Chaisemartin, Clément, and Xavier D'Haultfœuille. 2020. "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review* 110 (9): 2964–2996.

Charles, Kerwin Kofi, Erik Hurst, and Mariel Schwartz. 2019. "The Transformation of Manufacturing and the Decline in US Employment." *NBER Macroeconomics Annual* 33 307–372.

Chetty, Raj, John N. Friedman, and Jonah E. Rockoff. 2014. "Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates." *American Economic Review* 104 (9): 2593–2632.

Chetty, Raj, and Nathaniel Hendren. 2018a. "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects." *The Quarterly Journal of Economics* 133 (3): 1107–1162.

Chetty, Raj, and Nathaniel Hendren. 2018b. "The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates." *The Quarterly Journal of Economics* 133 (3): 1163–1228.

Ciesielski, Thomas, Reethi Iyengar, Amit Bothra, David Tomala, Geoffrey Cislo, and Brian F. Gage. 2017. "A Tool to Assess Risk of De Novo Opioid Abuse or Dependence." *The American Journal of Medicine* 129 (7): 699–705.

Corredor-Waldron, Adriana, and Janet Currie. 2022. "Tackling the Substance Use Disorder Crisis: The Role of Access to Treatment Facilities." *Journal of Health Economics* 81 102579.

Correia, Sergio, Paulo Guimarães, and Tom Zylkin. 2020. "Fast Poisson estimation with high-dimensional fixed effects." *The Stata Journal: Promoting communications on statistics and Stata* 20 (1): 95–115.

Cubanski, Juliette, Tricia Neuman, and Anthony Damico. 2016. *Medicare's Role for People Under Age 65 with Disabilities*. Kaiser Family Foundation Report.

Currie, Janet, Jonas Jin, and Molly Schnell. 2019. "US Employment and Opioids: Is There a Connection?." In *Health and Labor Markets*, 253–280, Emerald Publishing Limited.

Currie, Janet, and Hannes Schwandt. 2021. "The Opioid Epidemic Was Not Caused by Economic Distress but by Factors That Could Be More Rapidly Addressed." *The ANNALS of the American Academy of Political and Social Science* 695 (1): 276–291.

Cutler, David M., and Edward L. Glaeser. 2021. "When Innovation Goes Wrong: Technological Regress and the Opioid Epidemic." *Journal of Economic Perspectives* 35 (4): 171–96.

Dean, Adam, and Simeon Kimmel. 2019. "Free Trade and Opioid Overdose Death in the United States." *SSM - Population Health* 8 100409.

Deyo, Richard A., Sara E. Hallvik, Christi Hildebran et al. 2017. "Association Between Initial Opioid Prescribing Patterns and Subsequent Long-Term Use Among Opioid-Naïve Patients: A Statewide Retrospective Cohort Study." *Journal of General Internal Medicine* 32 (1): 21–27.

Doleac, Jennifer L., and Anita Mukherjee. 2018. "The Moral Hazard of Lifesaving Innovations: Naloxone Access, Opioid Abuse, and Crime." *IZA Discussion Paper No. 11489*.

Dunn, Kate M., Kathleen W. Saunders, Carolyn M. Rutter et al. 2010. "Overdose and Prescribed Opioids: Associations Among Chronic Non-cancer Pain Patients." *Annals of Internal Medicine* 152 (2): 85–92.

Edlund, Mark J., Bradley C. Martin, Joan E. Russo, Andrea DeVries, Jennifer B. Braden, and Mark D. Sullivan. 2014. "The Role of Opioid Prescription in Incident Opioid Abuse and Dependence Among Individuals With Chronic Noncancer Pain: The Role of Opioid Prescription." *The Clinical Journal of Pain* 30 (7): 557–564.

Eichmeyer, Sarah, and Jonathan Zhang. 2022. "Pathways into Opioid Dependence: Evidence from Practice Variation in Emergency Departments." *American Economic Journal: Applied Economics* 14 (4): 271–300.

Evans, William N., Ethan M. J. Lieber, and Patrick Power. 2019. "How the Reformulation of OxyContin Ignited the Heroin Epidemic." *The Review of Economics and Statistics* 101 (1): 1–15.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams. 2016. "Sources of Geographic Variation in Health Care: Evidence From Patient Migration." *The Quarterly Journal of Economics* 131 (4): 1681–1726.

Fischer, Benedikt, Anna Lusted, Michael Roerecke, Benjamin Taylor, and Jürgen Rehm. 2012. "The Prevalence of Mental Health and Pain Symptoms in General Population Samples Reporting Nonmedical Use of Prescription Opioids: A Systematic Review and Meta-Analysis." *The Journal of Pain* 13 (11): 1029–1044.

Guy, Gery P., Kun Zhang, Michele K. Bohm, Jan Losby, Brian Lewis, Randall Young, Louise B. Murphy, and Deborah Dowell. 2017. "Vital Signs: Changes in Opioid Prescribing in the United States, 2006-2015." *MMWR. Morbidity and Mortality Weekly Report* 66.

Hall, Aron J., Joseph E. Logan, Robin L. Toblin, James A. Kaplan, James C. Kraner, Danae Bixler, Alex E. Crosby, and Leonard J. Paulozzi. 2008. "Patterns of Abuse Among Unintentional Pharmaceutical Overdose Fatalities." *JAMA* 300 (22): 2613–2620.

Harris, Matthew C., Lawrence M. Kessler, Matthew N. Murray, and Beth Glenn. 2020. "Prescription Opioids and Labor Market Pains: The Effect of Schedule II Opioids on Labor Force Participation and Unemployment." *Journal of Human Resources* 55 (4): 1319–1364.

Hausman, J.A., Jason Abrevaya, and F.M. Scott-Morton. 1998. "Misclassification of the dependent variable in a discrete-response setting." *Journal of Econometrics* 87 (2): 239–269.

Heckman, James J., Hidehiko Ichimura, and Petra E. Todd. 1997. "Matching As An Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme." *The Review of Economic Studies* 64 (4): 605–654.

Hill, Maureen V., Michelle L. McMahon, Ryland S. Stucke, and Richard J. Barth. 2017. "Wide Variation and Excessive Dosage of Opioid Prescriptions for Common General Surgical Procedures." *Annals of Surgery*.

Hollingsworth, Alex, Christopher J. Ruhm, and Kosali Simon. 2017. "Macroeconomic Conditions and Opioid Abuse." *Journal of Health Economics* 56 222 – 233.

Ives, Timothy J., Paul R. Chelminski, Catherine A. Hammett-Stabler et al. 2006. "Predictors of Opioid Misuse in Patients with Chronic Pain: A Prospective Cohort Study." *BMC Health Services Research* 6 46.

Jeffery, Molly M., W. Michael Hooten, Erik P. Hess et al. 2018. "Opioid Prescribing for Opioid-Naive Patients in Emergency Departments and Other Settings: Characteristics of Prescriptions and Association With Long-Term Use." *Annals of Emergency Medicine* 71 (3): 326–336.e19.

Jena, Anupam B., Dana Goldman, Lesley Weaver, and Pinar Karaca-Mandic. 2014. "Opioid Prescribing by Multiple Providers in Medicare: Retrospective Observational Study of Insurance Claims." *BMJ* 348 g1393.

Kaestner, Robert, and Engy Ziedan. 2023. "Effects of prescription opioids on employment, earnings, marriage, disability and mortality: Evidence from state opioid control policies." *Labour Economics* 82 102358.

Katz, Josh. 2017. "Short Answers to Hard Questions About the Opioid Crisis." *The New York Times*.

Khan, Nazleen F., Brian T. Bateman, Joan E. Landon, and Joshua J. Gagne. 2019. "Association of Opioid Overdose With Opioid Prescriptions to Family Members." *JAMA Internal Medicine* 179 (9): 1186–1192.

Kiang, Mathew V, Keith Humphreys, Mark R Cullen, and Sanjay Basu. 2020. "Opioid prescribing patterns among medical providers in the United States, 2003-17: retrospective, observational study." *BMJ* l6968.

Kilby, Angela E. 2015. "Opioids for the Masses: Welfare Tradeoffs in the Regulation of Narcotic Pain Medications." `http://economics.mit.edu/files/11150`, accessed on 2020-11-15.

Klimas, Jan, Lauren Gorfinkel, Nadia Fairbairn, Laura Amato, Keith Ahamad, Seonaid Nolan, David L. Simel, and Evan Wood. 2019. "Strategies to Identify Patient Risks of Prescription Opioid Addiction When Initiating Opioids for Pain: A Systematic Review." *JAMA Network Open* 2 (5): e193365–e193365.

Kochanek, Kenneth D., Sherry L. Murphy, Jiaquan Xu, and Elizabeth Arias. 2019. "Deaths: Final Data for 2017." *National Vital Statistics Reports* 68 (9): 1–76.

Krueger, Alan B. 2017. "Where Have All the Workers Gone? An Inquiry into the Decline of the U.S. Labor Force Participation Rate." *Brookings Papers on Economic Activity* 2017 (2): 1–87.

Kuo, Yong-Fang, Mukaila A. Raji, and James S. Goodwin. 2019. "Association of Disability With Mortality From Opioid Overdose Among US Medicare Adults." *JAMA Network Open* 2 (11): e1915638.

Laird, Jessica, and Torben Nielsen. 2016. "The Effects of Physician Prescribing Behaviors on Prescription Drug Use and Labor Supply: Evidence from Movers in Denmark." *Working Paper*, accessed on 2021-03-31.

Larrick, Amy. 2014. "Request for Comments: Enhancements to the Star Ratings for 2016 and Beyond." `https://www.cms.gov/Medicare/Prescription-Drug-Coverage/PrescriptionDrugCovGenIn/Downloads/2016-Request-for-Comments-v-11_25_2014.pdf`.

Logan, Joseph, Ying Liu, Leonard Paulozzi, Kun Zhang, and Christopher Jones. 2013. "Opioid Prescribing in Emergency Departments: The Prevalence of Potentially Inappropriate Prescribing and Misuse." *Medical Care* 51 (8): 646–653.

Lyapustina, Tatyana, Lainie Rutkow, Hsien-Yen Chang, Matthew Daubresse, Alim F. Ramji, Mark Faul, Elizabeth A. Stuart, and G. Caleb Alexander. 2016. "Effect of a Pill Mill Law on Opioid Prescribing and Utilization: The Case of Texas." *Drug and Alcohol Dependence* 159 190–197.

Maclean, Johanna Catherine, Justine Mallatt, Christopher J. Ruhm, and Kosali Simon. 2022. "The Opioid Crisis, Health, Healthcare, and Crime: A Review of Quasi-Experimental Economic Studies." *The ANNALS of the American Academy of Political and Social Science* 703 (1): 15–49.

McDonald, Douglas C., Kenneth Carlson, and David Izrael. 2012. "Geographic Variation in Opioid Prescribing in the U.S.." *The Journal of Pain* 13 (10): 988–996.

Meara, Ellen, Jill R. Horwitz, Wilson Powell, Lynn McClelland, Weiping Zhou, A. James O'Malley, and Nancy E. Morden. 2016. "State Legal Restrictions and Prescription-Opioid Use among Disabled Adults." *New England Journal of Medicine* 375 (1): 44–53.

Molitor, David. 2018. "The Evolution of Physician Practice Styles: Evidence from Cardiologist Migration." *American Economic Journal: Economic Policy* 10 (1): 326–356.

Morden, Nancy E., Jeffrey C. Munson, Carrie H. Colla, Jonathan S. Skinner, Julie P. W. Bynum, Weiping Zhou, and Ellen Meara. 2014. "Prescription Opioid Use Among Disabled Medicare Beneficiaries: Intensity, Trends, and Regional Variation." *Medical Care* 52 (9): 852–859.

Moré, Jorge J. 1978. "The Levenberg-Marquardt algorithm: Implementation and Theory." In *Lecture Notes in Mathematics*, 105–116, Springer Berlin Heidelberg.

Mulligan, Casey B. 2024. "Prices and Policies in Opioid Markets." *Journal of Political Economy* 132 (10): 3461–3499.

National Center for Health Statistics. 2021. *CDC WONDER: National Drug Overdose (OD) Deaths, 1999-2019*. `https://wonder.cdc.gov/mcd.html`.

Packham, Analisa. 2022. "Syringe exchange programs and harm reduction: New evidence in the wake of the opioid epidemic." *Journal of Public Economics* 215 104733.

Parente, Stephen T., Susan S. Kim, Michael D. Finch, Lisa A. Schloff, Thomas S. Rector, Raafat Seifeldin, and J. David Haddox. 2004. "Identifying Controlled Substance Patterns of Utilization Requiring Evaluation Using Administrative Claims Data." *The American Journal of Managed Care* 10 (11 Pt 1): 783–790.

Paulozzi, Leonard J., Edwin M. Kilbourne, Nina G. Shah, Kurt B. Nolte, Hema A. Desai, Michael G. Landen, William Harvey, and Larry D. Loring. 2012. "A History of Being Prescribed Controlled Substances and Risk of Drug Overdose Death." *Pain Medicine* 13 (1): 87–95.

Powell, David, Rosalie Liccardo Pacula, and Erin Taylor. 2020. "How Increasing Medical Access to Opioids Contributes to the Opioid Epidemic: Evidence from Medicare Part D." *Journal of Health Economics* 71 102286.

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

Rees, Daniel I., Joseph J. Sabia, Laura M. Argys, Dhaval Dave, and Joshua Latshaw. 2019. "With a Little Help from My Friends: The Effects of Good Samaritan and Naloxone Access Laws on Opioid-related Deaths." *The Journal of Law and Economics* 62 (1): 1–27.

Rice, J. Bradford, Alan G. White, Howard G. Birnbaum, Matt Schiller, David A. Brown, and Carl L. Roland. 2012. "A Model to Identify Patients at Risk for Prescription Opioid Abuse, Dependence, and Misuse." *Pain Medicine* 13 (9): 1162–1173.

Rikard, S. Michaela. 2023. "Chronic Pain Among Adults — United States, 2019–2021." *MMWR. Morbidity and Mortality Weekly Report* 72.

Rough, Kathryn, Krista F. Huybrechts, Sonia Hernandez-Diaz, Rishi J. Desai, Elisabetta Patorno, and Brian T. Bateman. 2019. "Using Prescription Claims to Detect Aberrant Behaviors with Opioids: Comparison and Validation of 5 Algorithms." *Pharmacoepidemiology and Drug Safety* 28 (1): 62–69, 29687539[pmid].

Rubin, Donald B. 1981. "The Bayesian Bootstrap." *The Annals of Statistics* 9 (1): 130–134.

Ruhm, Christopher J. 2018. "Deaths of Despair or Drug Problems?" *NBER Working Paper No. 24188*.

Rutkow, Lainie, Hsien-Yen Chang, Matthew Daubresse, Daniel W. Webster, Elizabeth A. Stuart, and G. Caleb Alexander. 2015. "Effect of Florida's Prescription Drug Monitoring Program and Pill Mill Laws on Opioid Prescribing and Use." *JAMA Internal Medicine* 175 (10): 1642–1649.

Sacks, Daniel W., Alex Hollingsworth, Thuy Nguyen, and Kosali Simon. 2021. "Can policy affect initiation of addictive substance use? Evidence from opioid prescribing." *Journal of Health Economics* 76 102397.

Schieber, Lyna Z., Jr Guy, Gery P., Puja Seth, Randall Young, Christine L. Mattson, Christina A. Mikosz, and Richard A. Schieber. 2019. "Trends and Patterns of Geographic Variation in Opioid Prescribing Practices by State, United States, 2006-2017." *JAMA Network Open* 2 (3): e190665–e190665.

Schnell, Molly. 2022. "Physician Behavior in the Presence of a Secondary Market: The Case of Prescription Opioids." *Working Paper*.

Schnell, Molly, and Janet Currie. 2018. "Addressing the Opioid Epidemic: Is There a Role for Physician Education?." *American Journal of Health Economics* 4 (3): 383–410.

Scully, Rebecca E., Andrew J. Schoenfeld, Wei Jiang et al. 2018. "Defining Optimal Length of Opioid Pain Medication Prescription After Common Surgical Procedures." *JAMA Surgery*.

Shah, Anuj, Corey J. Hayes, and Bradley C. Martin. 2017. "Factors Influencing Long-Term Opioid Use Among Opioid Naive Patients: An Examination of Initial Prescription Characteristics and Pain Etiologies." *The Journal of Pain* 18 (11): 1374–1383.

Skinner, Jonathan. 2011. "Chapter Two - Causes and Consequences of Regional Variations in Health." In *Volume 2 of Handbook of Health Economics*, edited by Pauly, Mark V., Thomas G. Mcguire, and Pedro P. Barros 45–93, Elsevier.

Social Security Administration. 2017. *Annual Statistical Supplement to the Social Security Bulletin, 2016*.

Song, Yunjie, Jonathan Skinner, Julie Bynum, Jason Sutherland, John E. Wennberg, and Elliott S. Fisher. 2010. "Regional Variations in Diagnostic Practices." *New England Journal of Medicine* 363 (1): 45–53.

Stone, Elizabeth M., Lainie Rutkow, Mark C. Bicket, Colleen L. Barry, G. Caleb Alexander, and Emma E. McGinty. 2020. "Implementation and Enforcement of State Opioid Prescribing Laws." *Drug and alcohol dependence* 213 108107.

Sullivan, Mark D., Mark J. Edlund, Ming-Yu Fan, Andrea Devries, Jennifer Brennan Braden, and Bradley C. Martin. 2010. "Risks for Possible and Probable Opioid Misuse Among Recipients of Chronic Opioid Therapy in Commercial and Medicaid Insurance Plans: The TROUP Study." *Pain* 150 (2): 332–339.

Sun, Eric C., Beth D. Darnall, Laurence C. Baker, and Sean Mackey. 2016. "Incidence of and Risk Factors for Chronic Opioid Use Among Opioid-Naive Patients in the Postoperative Period." *JAMA Internal Medicine* 176 (9): 1286–1293.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." *Journal of Econometrics* 225 (2): 175–199, Themed Issue: Treatment Effect 1.

Turk, Dennis C., Kimberly S. Swanson, and Robert J. Gatchel. 2008. "Predicting Opioid Misuse by Chronic Pain Patients: A Systematic Review and Literature Synthesis." *The Clinical Journal of Pain* 24 (6): 497–508.

United States General Accounting Office. 1991. *The Crack Cocaine Epidemic: Health Consequences and Treatment*.

Venkataramani, Atheendar S., Elizabeth F. Bair, Rourke L. O'Brien, and Alexander C. Tsai. 2020. "Association Between Automotive Assembly Plant Closures and Opioid Overdose Mortality in the United States: A Difference-in-Differences Analysis." *JAMA Internal Medicine* 180 (2): 254–262.

Webster, Lynn R. 2017. "Risk Factors for Opioid-Use Disorder and Overdose." *Anesthesia & Analgesia* 125 (5): 1741–1748.

Wei, Yu-Jung Jenny, Cheng Chen, Amir Sarayani, and Almut G. Winterstein. 2019. "Performance of the Centers for Medicare & Medicaid Services' Opioid Overutilization Criteria for Classifying Opioid Use Disorder or Overdose." *JAMA* 321 (6): 609–611.

White, Alan G., Howard G. Birnbaum, Matt Schiller, Jackson Tang, and Nathaniel P. Katz. 2009. "Analytic Models to Identify Patients at Risk for Prescription Opioid Abuse." *The American Journal of Managed Care* 15 (12): 897–906.

Wooldridge, Jeffrey M. 2021. "Two-Way Fixed Effects, the Two-Way Mundlak Regression, and Difference-in-Differences Estimators." *Available at SSRN: http://dx.doi.org/10.2139/ssrn.3906345*.

Yang, Zhuo, Barth Wilsey, Michele Bohm, Meghan Weyrich, Kakoli Roy, Dominique Ritley, Christopher Jones, and Joy Melnikow. 2015. "Defining Risk of Prescription Opioid Overdose: Pharmacy Shopping and Overlapping Prescriptions Among Long-term Opioid Users in Medicaid." *The Journal of Pain* 16 (5): 445–453.

Figure 1: Geographic Variation in Rates of Risky Use

(a) Map of Average Risky Use Rates

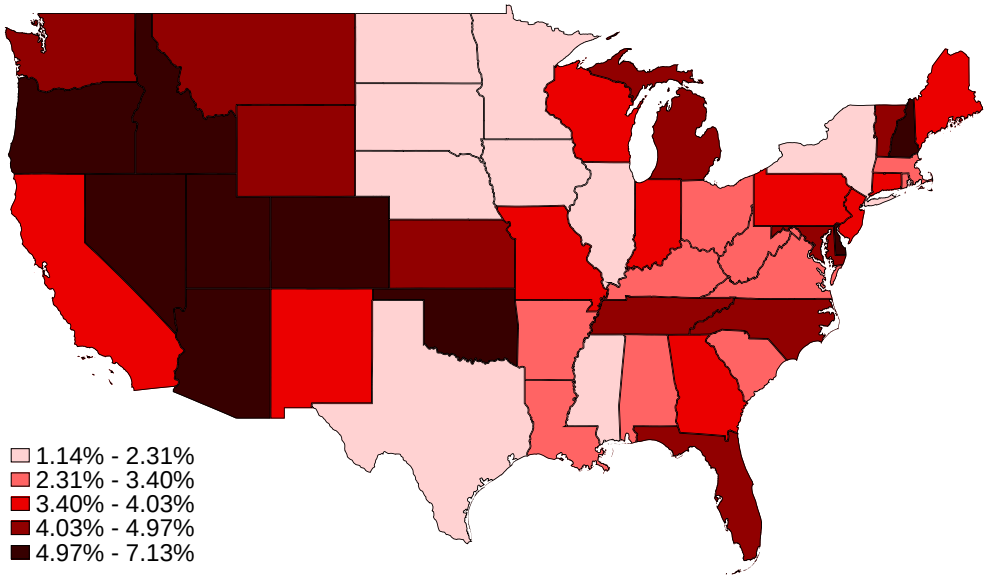

(b) Risky Use by Year

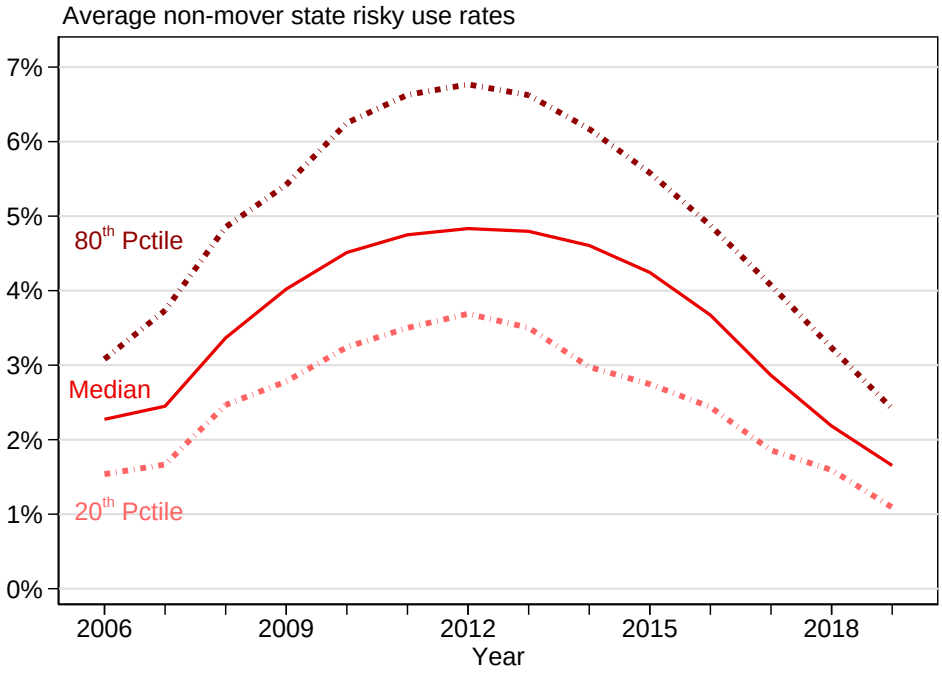

Notes: Figure reports state-level averages for the rate of risky opioid use among all non-mover years (N = 2,755,447 enrollee-years). Panel A presents a map of the state-level averages over the full sample period (2006 - 2019) while Panel B presents the 20th, 50th, and 80th percentiles of risky use among states in each calendar year.

Figure 2: Interpreting Treatment Effects in Event Study Specification

(a) Case 1: Only Person Effects Vary

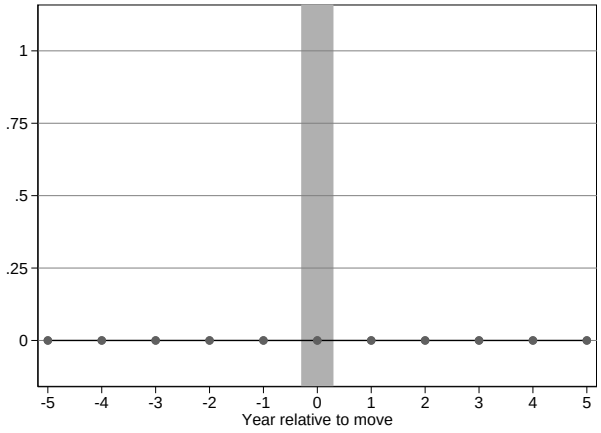

(b) Case 2: Only Availability Effects Vary

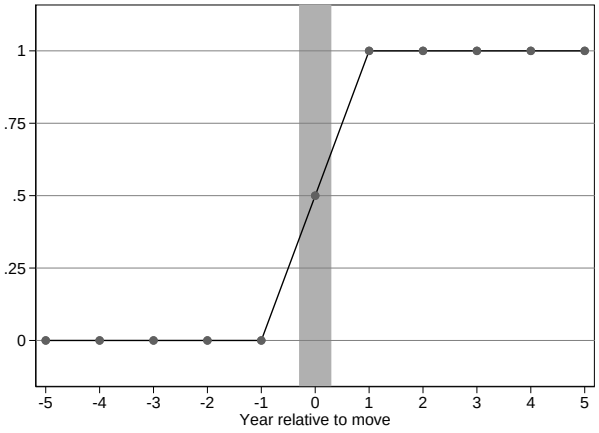

(c) Case 3: Both Person and Availability Effects Vary

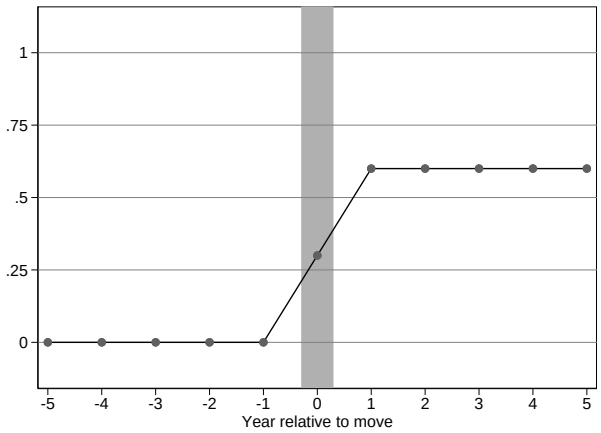

(d) Case 4: Person, Availability, and Addiction Effects Vary

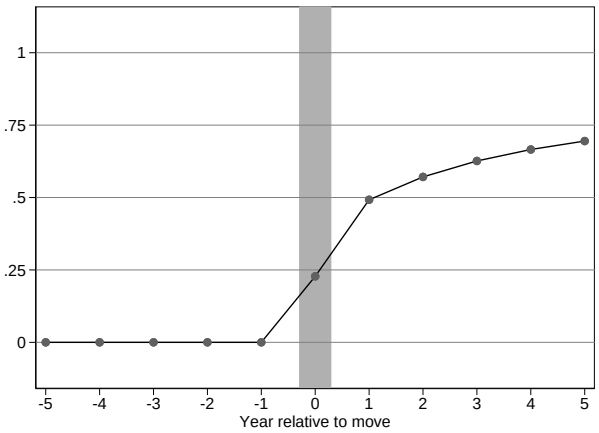

Notes: Figure shows the patterns of event study coefficients μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—as calculated according to equation (5) in a theoretical setting with a single cohort moving between two locations and the coefficients plotted on the y-axis. The cohort is randomly sampled from the origin, and risky use rates are determined according to the parameters as described in Section 3. We choose an illustrative set of parameters for each example. In addition, we assume that in relative year 0, there is a 50% chance that each individual is in their origin location and a 50% chance that they are in their destination location. In Panel A, only person-based factors differ between the two locations. In Panel B, only availability effects differ between the two locations. In Panel C, both person-based factors and availability effects differ between the two locations, and we choose a case where the ratio of shares addicted and availability effects in the origin and destination are exactly constant. Finally, in Panel D, the two locations differ in their person-based factors, availability effects, and their place-based effects on transitions to addiction. The parameters used in Panel C are $\pi^{+}_{o(c)} = \pi^{+}_{d(c)} = 0.01$, $\pi^{-}_{o(c)} = \pi^{-}_{d(c)} = 0.16$, with $\gamma_{o(c)r} = 0.10$, $\gamma_{d(c)r} = 0.40$, $a_{o(c)r} = 0.10$ and $a_{d(c)r} = 0.15$ for all $r \in [-5,5]$, and $\eta^{+}_{i} = 0.02$, $\eta^{-}_{i} = 0.01$for all $i \in \mathscr{J}_{o(c)}$ while $\eta^{+}_{i} = 0.01$, $\eta^{-}_{i} = 0.02$ for all $i \in \mathscr{J}_{d(c)}$. The parameters used in Panel D are $\pi^{+}_{o(c)} = 0.05$, $\pi^{+}_{d(c)} = 0.03$, $\pi^{-}_{o(c)} = 0.03$, $\pi^{-}_{d(c)} = 0.10$, with $\gamma_{o(c)r} = 0.30$ and $\gamma_{d(c)r} = 0.26$ for $r \in [-5,5]$, $a_{o(c),-5} = a_{d(c),-5} = 0.10$, and $\eta^{+}_{i} = 0.05$, $\eta^{-}_{i} = 0.02$ for all $i \in \mathscr{J}_{o(c)}$ while $\eta^{+}_{i} = 0.04$, $\eta^{-}_{i} = 0.05$ for all $i \in \mathscr{J}_{d(c)}$.

Figure 3: Distribution of Origin - Destination Differences in Rate of Risky Use

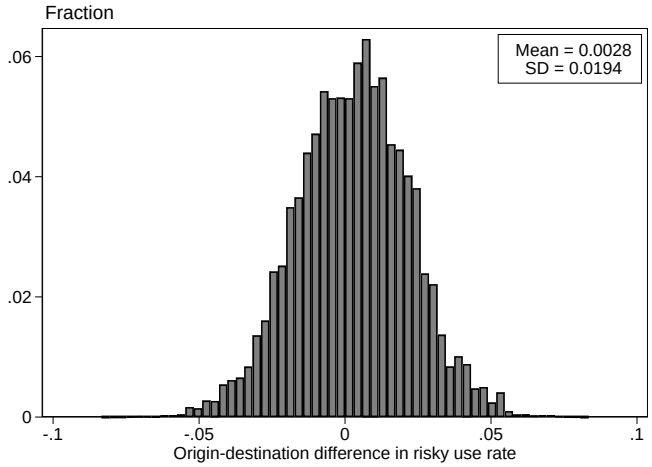

Notes: Figure shows the distribution across movers of the difference in the average rates of risky opioid use between their origin and destination states in a given year. The sample is all mover-years ($N = 726,146$ mover-years).

Figure 4: Event Study

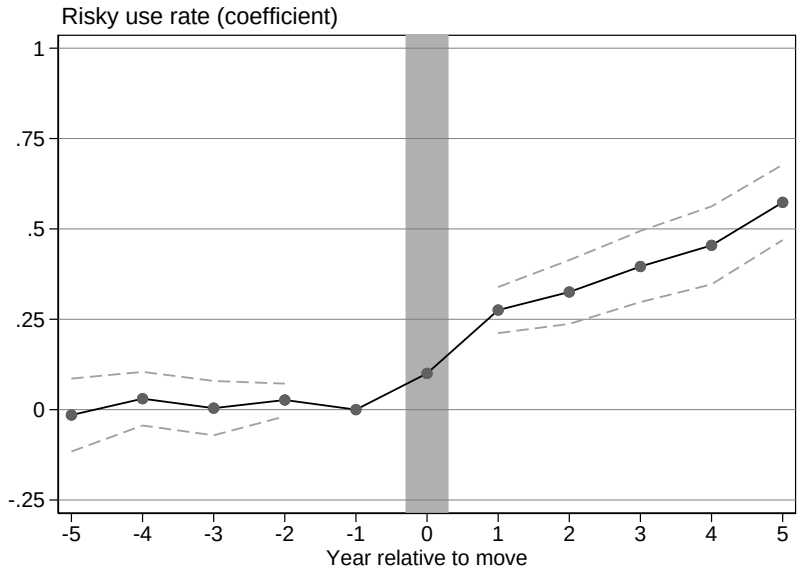

Notes: Figure shows estimates from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort. The dashed lines are upper and lower bounds of the 95% confidence interval. Confidence intervals are computed using 50 iterations of the Bayesian bootstrap. Note that our definition of $\hat{T}_{cr}$ in Section 4.1 implies that the coefficient in relative year -1 is zero by construction. The sample is all mover-years ($N = 726,146$ mover-years).

Figure 5: Change in Rate of Risky Use by Size of Move

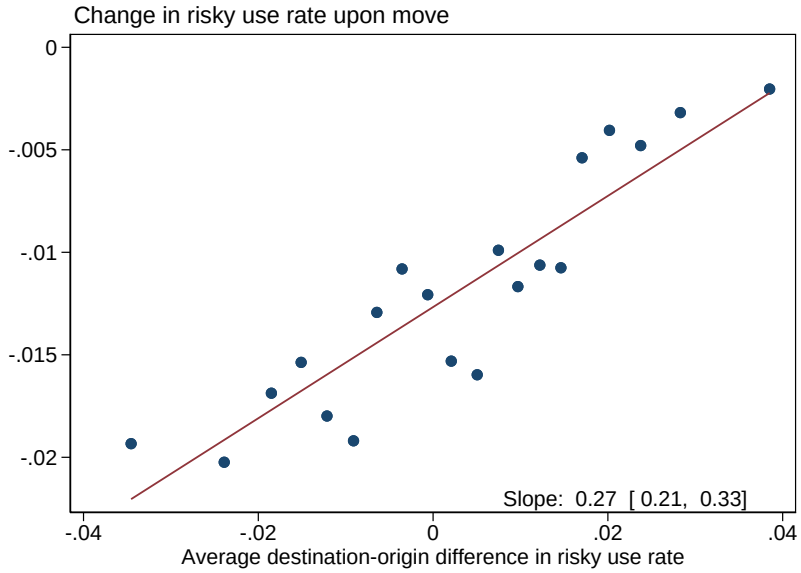

Notes: Figure shows changes in the rate of risky use from before to after move as a function of the size of the move ($\hat{\delta}_{cr}$). For each mover, we calculate the average difference $\hat{\delta}_{cr}$ in the rate of risky use between their origin and destination states for all years they are observed between one and five years pre-move and one and five years post-move. We then group these average differences—$\hat{\delta}_{cr}$—into ventiles. The x-axis displays the mean of $\hat{\delta}_{cr}$ for movers in each ventile. The y-axis shows, for each ventile, the average risky use rate one to five years post-move minus the average risky use rate one to five years pre-move, averaged within the ventiles. The line of best fit is obtained from a simple OLS regression using the 20 data points corresponding to movers, and its slope is reported in the graph with a 95% confidence interval in brackets. The sample is all mover-years between one and five years pre-move and one and five years post-move ($N = 680,617$ mover-years).

Figure 6: Event Studies - Naives and Prior Users

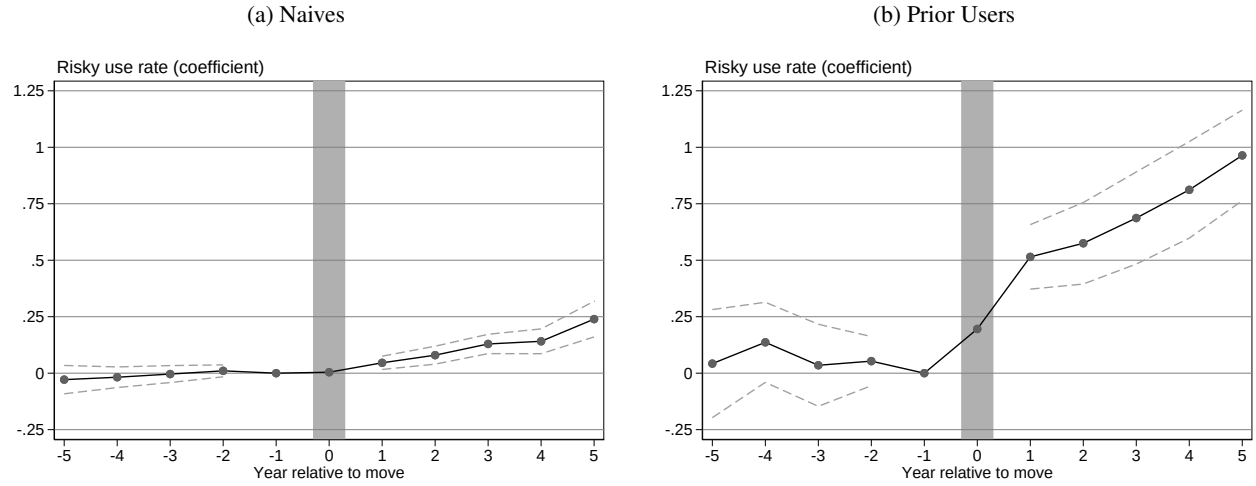

Notes: Figure shows estimates from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort. We show these estimates separately for opioid naive ("naive") and prior users. "Naives" are those with no opioid use in relative year -1, while prior users filled at least one opioid prescription in that year. For each subsample, we extend the vector of observables used for matching to include opioid use in the calendar year corresponding to the year before moving. The dashed lines are upper and lower bounds of the 95% confidence interval. Confidence intervals are computed using 50 iterations of the Bayesian bootstrap. Note that our definition of $\hat{T}_{cr}$ in Section 4.1 implies that the coefficient in relative year -1 is zero by construction. The samples are 412,821 mover-years (naives) and 312,866 mover-years (prior users).

Figure 7: Event Study in Simulation vs. Data

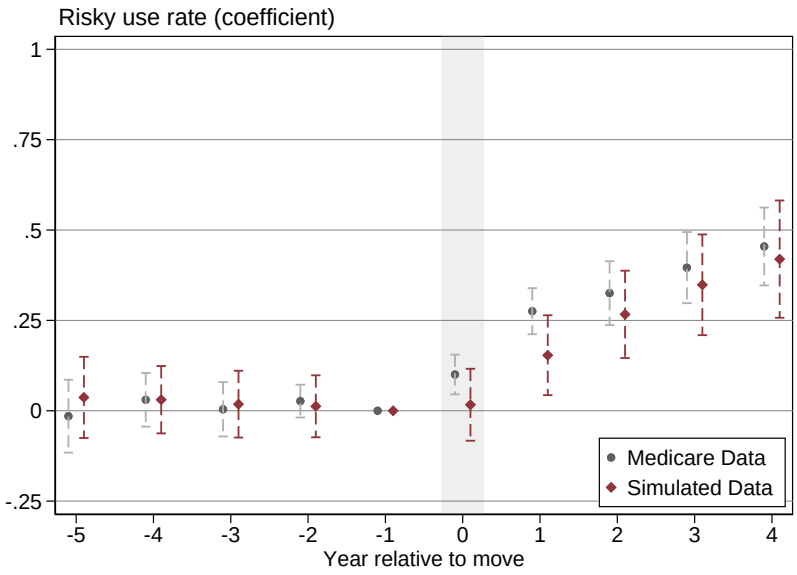

Notes: Figure shows estimates from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—as estimated on simulated data and the Medicare data, where the weights are given by the number of movers observed in each cohort. The simulated data is constructed at the enrollee-year level by using our estimated parameters to simulate individual-level addiction and risky use outcomes for enrollee-years in the baseline sample. Red diamonds show the aggregate coefficients estimated from our simulated event study, where the yearly rates of risky opioid use among movers and non-movers are simulated according to the model parameterizations described in Section 5 and confidence intervals are calculated across fifty simulations. Gray dots and dashed lines are identical to Figure (4) and plot our aggregate event study coefficients and standard errors as estimated on our baseline sample. Our definition of $\hat{T}_{cr}$ in Section 4.1 implies that the coefficient in relative year -1 is zero by construction.

Figure 8: Geographic Variation in Calibrated Parameters

(a) Availability (γ_j)

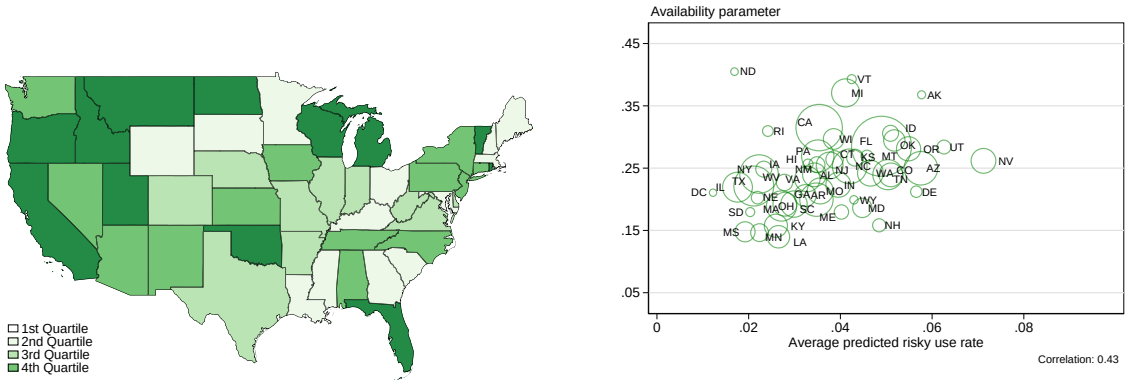

(b) Person-Based Addiction Transitions (η_i)

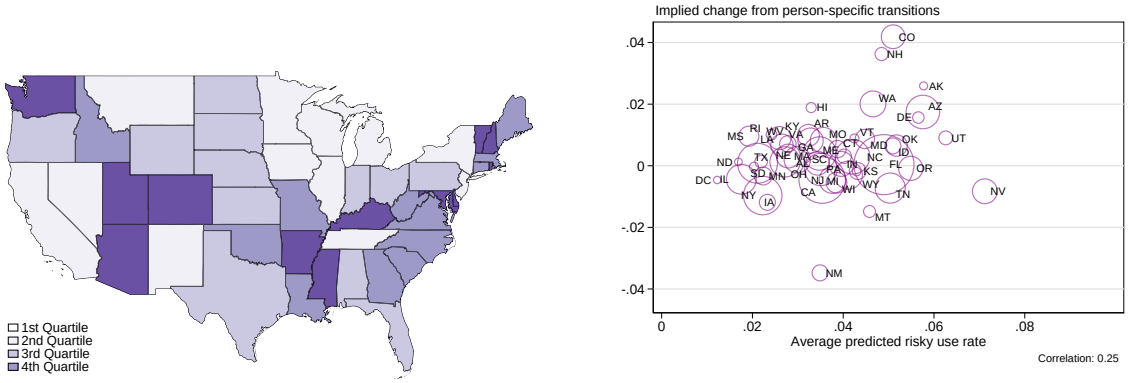

(c) Place-Based Addiction Transitions (π_j)

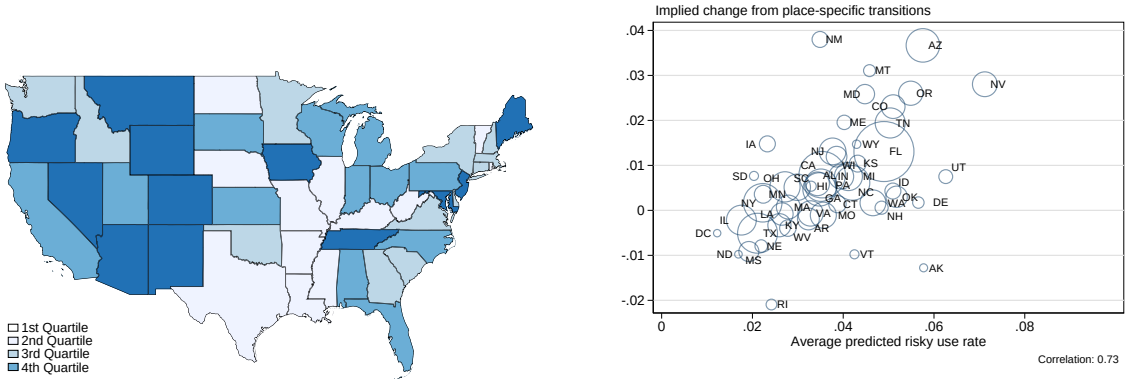

Notes: Figure presents three panels with both maps and scatterplots of parameter estimates for various channels in our model. Panel A presents the place-specific availability parameter, while Panels B and C present person-specific and place-specific factors affecting addiction transitions respectively. The weighted correlation coefficient is presented in the bottom right corner, where weights are given by the number of mover-year observations in the estimation sample where a mover was observed in the state. Markers are scaled accordingly. The magnitude of addiction transition parameters is measured by the change in shares addicted that would result after one year if all states shared an initial share of 0.10 and differential addiction transitions occurred only according to variation in the relevant set of addiction transition parameters. The implied one-year change from place/person specific addiction transitions specifically computes $\Delta\bar{a} = \bar{a}_0(1-\pi^- - \eta^-) + (1-\bar{a}_0)(\pi^+ + \eta^+) - \bar{a}_0$, where $\bar{a}_0$ is set to the global share of addiction in our model estimation and the other set of parameters (place vs. person) are held to their median values. We discuss the calibration of a global share of addiction further in Appendix E.

Figure 9: CZ-Level Correlates of Risky Use and Model Parameters

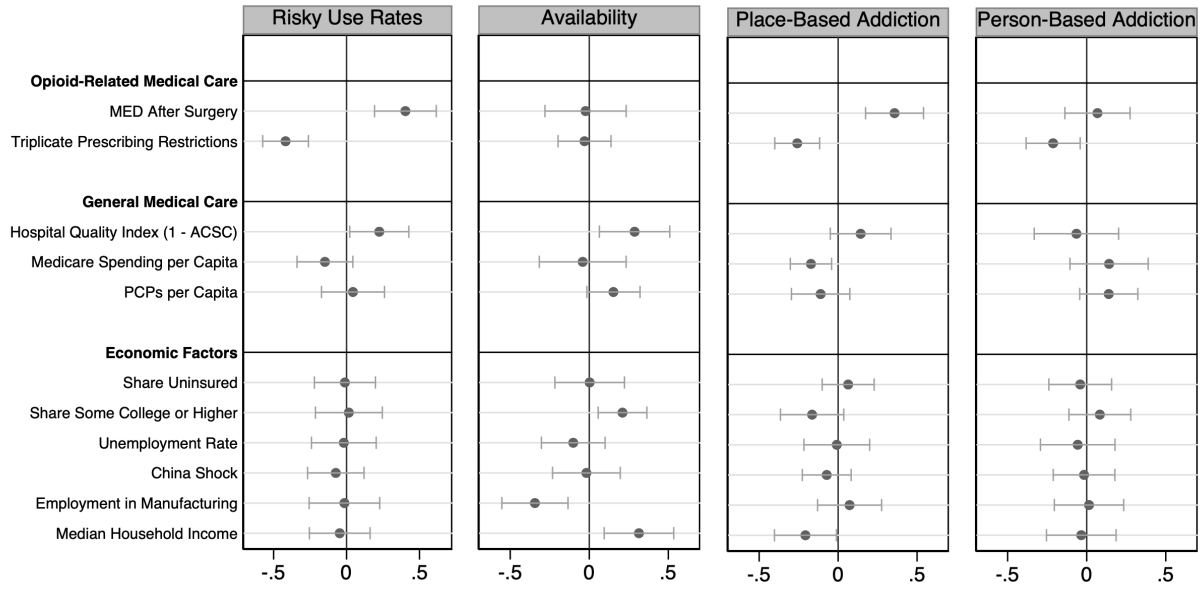

Notes: Figure presents coefficients from bivariate regressions using commuting-zone level average characteristics on our model parameters. Horizontal bars show the 95 percent confidence intervals, calculated with heteroskedasticity-robust standard errors. All outcomes and covariates are normalized to be in units of standard deviations. The first panel from the left presents coefficients using the average risky use of the commuting zone in our sample as the outcome. The second panel from the left presents coefficients using the estimated availability parameter ($\hat{\gamma}_j$) of that commuting zone. The third panel from the left presents the steady state share of addiction in the commuting zone implied by using variation in the estimated place-based addiction parameters ($\hat{\pi}_j$) while holding estimated person-based addiction parameters to their median value among all CZs. Finally, the right-most panel presents the steady state share of addiction in the commuting zone implied by using variation in the estimated person-based addiction parameters ($\hat{\eta}_i$) among locations while holding estimated place-based addiction parameters to their median value among all CZs. All demographic and economic outcomes are measured in the 2000 Census. "MED After Surgery" refers to the average morphine equivalent dose (MED) of the prescriptions patients filled in the two weeks following a set of common surgical procedures among individuals with traditional fee for service coverage in our sample. We measure hospital quality through the "ACSC Rate", or the rate of hospitalizations for ambulatory-care sensitive conditions (ACSC). High ACSC rates are considered a measure of poor healthcare quality, and so we define quality by the rate of avoided hospitalizations for these conditions (1 - ACSC). We present detailed definitions and data sources for these measures in Appendix H. The estimates shown in this figure (along with their standard errors) are reported in Appendix Table A.4.

Figure 10: The Time Path of Counterfactual Reductions on Risky Use

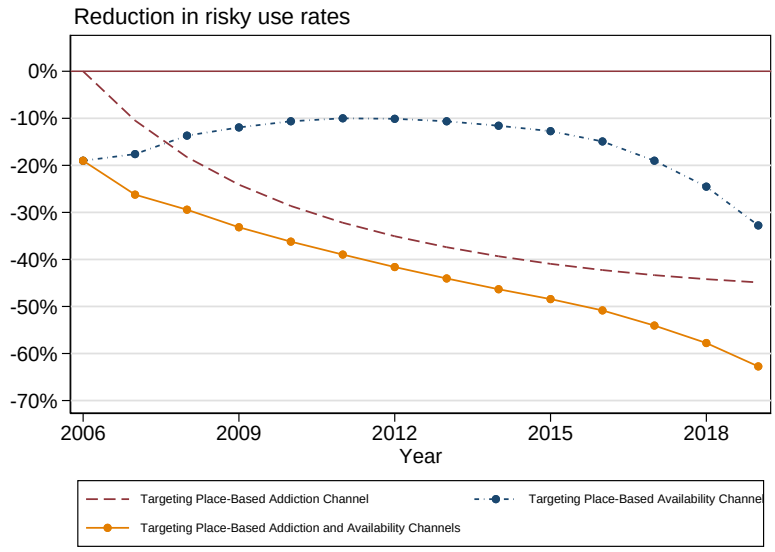

Notes: Figure shows the yearly effect of counterfactual policies that lower place-based addiction and place-based availability parameters. Each counterfactual policy is to reduce parameters in the relevant channel(s) by a standard deviation except for the probabilities of transitioning out of addiction, which we increase by a standard deviation. The average yearly risky use rate is computed by taking a weighted average across simulated risky use outcomes in each state-year for our non-mover sample, weighting by the number of non-movers in each state-year. Throughout, we maintain our standard parameter bounds of [0,1] by taking the minimum (maximum) of the reduced (increased) parameter and the bound as necessary. The estimates shown in this figure (along with their standard errors) are reported in Appendix Table A.6.

Figure 11: Analyzing the Impact of Pill Mill Policies

(a) Percentage Reduction in Risky Use Rates

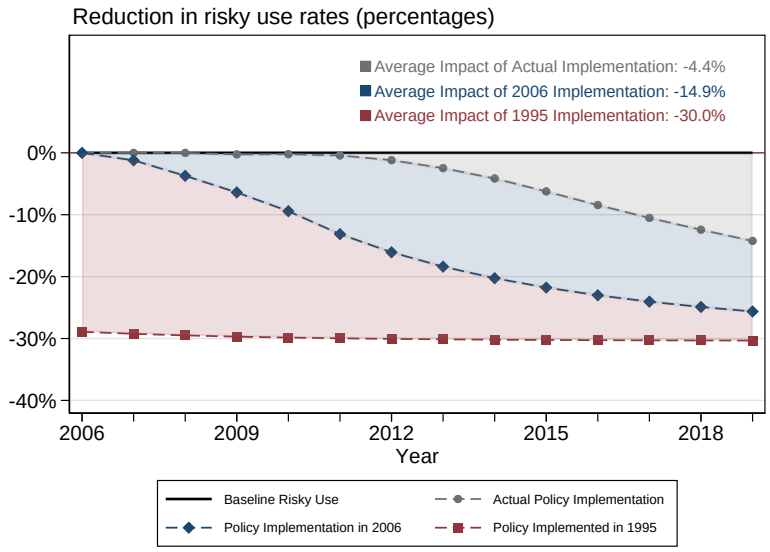

(b) Risky Use Rates

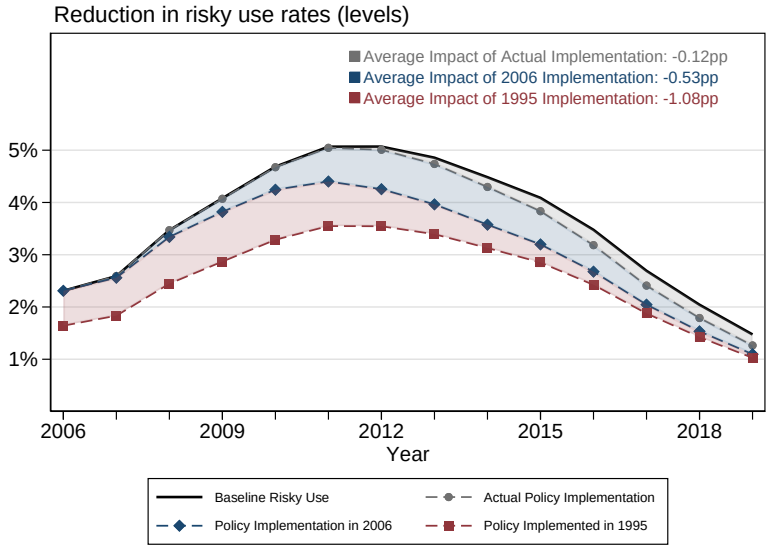

Notes: Figure shows the estimated simulated impact of pill mill laws in the 11 eventually-treated states under various counterfactual policy passage dates. These implementation counterfactuals maintain the gradual phase-in of policy effects over five years as described in Section 5.2. Panel A shows the reduction in the simulated risky use rates from 2006 to 2019 under the actual timing (gray), a 2006 implementation (blue) and a 1995 implementation (red). The impact is measured relative to the counterfactual average risky use rate implied by our model had they not been implemented. Panel B shows the simulated risky use rates directly. The average impact is estimated by averaging over calendar years.

Table 1: Summary Statistics

	(1) Movers	(2) Non-movers
Female	57%	52%
White	74%	71%
Black	19%	20%
Medicaid	57%	55%
Age:		
< 40	11%	10%
40 - 60	43%	41%
> 60	45%	49%
Average age	57.5	58.5
Region:		
Northeast	16%	19%
West	22%	18%
Midwest	18%	22%
South	43%	41%
Opioid use:		
Any opioids	41.1%	36.6%
Prescriptions in year before move ("prior user")	43.1%	
No prescriptions in year before move ("opioid naive")	56.9%	
Risky use	4.1%	3.5%
Number of enrollee-years	726,146	20,094,127
Number of enrollees	99,729	2,755,447

Notes: All rows except for the number of enrollee-years and enrollees report the share of enrollees or enrollee-years within the given population with the indicated characteristic. "Any Opioids" and "Risky Use" are averaged over all enrollee-years, while all other statistics are averaged at the enrollee-level, with "Region," "Medicaid," and "Age" measured in a reference year. This reference year is the year before move for movers, and a randomly assigned year for non-movers such that the distribution of reference years for non-movers mirrors that of movers.

Table 2: Event Study Coefficients - Rates of Risky Use

	(1)	(2)	(3)
	All	Naive	Prior User
1 year post-move	0.275	0.046	0.515
	(0.033)	(0.015)	(0.073)
5 years post-move	0.573	0.239	0.964
	(0.053)	(0.040)	(0.102)
Enrollees	99,729	57,334	42,341
Enrollee-years	726,146	412,821	312,866
Average Risky Use Rate 1 year pre-move	0.043	0.000	0.100

Notes: Table reports the coefficients and their bootstrapped standard errors (in parentheses) one and five years post-move as estimated from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort. Estimates are shown for the baseline sample ("All"), naive enrollees, and prior users, for movers observed in the year before move for whom we are able to observe matched non-movers. Opioid naive enrollees are those with no opioid use in relative year -1, while prior users filled at least one opioid prescription in that year. For the "naive" and "prior user" subsamples, we extend the vector of observables used for matching to include opioid use in the calendar year corresponding to the year before moving. Standard errors are computed using 50 iterations of the Bayesian bootstrap.

Table 3: Model-Estimated Policy Effects

	Policy Parameter	Estimates
	Full Policy Effects	
(1)	Addiction (ζ^{π}_{max})	-0.335
		(0.144)
(2)	Availability (ζ^{γ}_{max})	0.002
		(0.052)
	Timing Parameters	
(3)	Year of Implementation ($share_0$)	0.000
		(0.000)
(4)	First Year ($share_1$)	0.202
		(0.153)
(5)	Second Year ($share_2$)	0.468
		(0.056)
(6)	Third Year ($share_3$)	0.598
		(0.048)
(7)	Fourth Year ($share_4$)	0.764
		(0.048)

Notes: Table shows estimates of our policy-related model parameters. Rows (1) and (2) present the full effects of the policies, while rows (3) - (7) present the "share" of the policy in effect, reflecting the gradual implementation of the policy. Standard errors are computed using 50 iterations of the Bayesian bootstrap.

Table 4: The Effect of Counterfactual Reductions on Risky Use

	Predicted reduction in rate of risky use from a one standard deviation reduction in...	Reduction in risky use between 2006 - 2019
(1)	Person-based Parameters (η_i)	32.6%
		(4.8%)
(2)	Place-based Parameters (γ_j, π_j)	41.3%
		(4.9%)
	(a) Availability (γ_j)	13.9%
		(2.8%)
	(b) Place-based Addiction Transitions (π_j)	31.9%
		(5.2%)
(3)	All Addiction Transitions (π_j, η_i)	51.0%
		(8.2%)

Notes: Table reports the percent reduction in simulated overall risky use rates from 2006 to 2019 from a one standard deviation reduction in parameters. Each counterfactual policy is to reduce parameters in the relevant channel(s) by a standard deviation, except for the probabilities of transitioning out of addiction, which we increase by a standard deviation. We assume that the change in parameters occurs during 2006. The overall risky use rate is computed by taking a weighted average of simulated risky use over all state-years in our non-mover sample, weighting by the number of non-movers observed in each state-year. 95% confidence intervals are computed using 50 iterations of the Bayesian bootstrap. Throughout, we maintain our standard parameter bounds of [0,1] by taking the minimum (maximum) of the reduced (increased) parameter and the bound as necessary.

Table 5: The Effect of Model Addiction and Availability on Adverse Opioid-Related Outcomes

	Adverse Opioid-Related Outcomes			
	Risky Use	**Opioid-Related Poisoning**	**Any Drug-Related Poisoning**	**Opioid-Use Disorder**
	(1)	(2)	(3)	(4)
Availability	0.523***	0.011**	0.013	-0.009
	(0.066)	(0.005)	(0.015)	(0.029)
Place-Based Addiction	1.222***	0.027***	0.062***	0.031
	(0.125)	(0.006)	(0.015)	(0.019)
Person-Based Addiction	0.724***	0.028***	0.059***	-0.019
	(0.0532)	(0.003)	(0.013)	(0.026)
N	91	91	91	91
Mean of Outcome in Sample	3.325%	0.142%	0.604%	0.419%
Correlation with Risky Use	1.000	0.729	0.636	0.153

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

Notes: Table shows weighted regression coefficients of our model parameters on various adverse opioid-related outcomes and their heteroskedasticity-robust standard errors. We standardize our model parameters so that units are in standard deviations. This exercise uses commuting-zone level adverse opioid-related outcomes and our model parameters as estimated on a subsample of commuting zones. More details on the commuting-zone level estimation are provided in Appendix C. Commuting zones are weighted by the number of non-movers observed in commuting zone sample. For columns (2) - (4), we calculate commuting-zone level rates among non-mover enrollee years covered by traditional Medicare whom we also observe in the following year. Availability denotes the estimated availability parameter ($\hat{\gamma}_j$). Place-based addiction refers to the steady state share of addiction in the commuting zone implied by using variation in the estimated place-based addiction parameters ($\hat{\pi}_j$) while holding the estimated person-based addiction parameters to their median value among all CZs. Person-based addiction refers to the steady state share of addiction in the commuting zone implied by using variation in the estimated person-based addiction parameters ($\hat{\eta}_i$) for each location while holding estimated place-based addiction parameters to their median value among all CZs.

Appendix A: Defining Risky Opioid Use

Measuring risky opioid use and adverse opioid events

As discussed in the main text, clinicians and medical researchers have not come to a consensus on a gold standard measure of risky opioid use from claims data. However, the literature uses several proxies for risky use based on apparent hazardous prescription patterns.

While the simplest measure of hazardous prescriptions is the number of opioid prescriptions a patient fills over a fixed time period (Rice et al. 2012), a better measure of hazardous prescription behavior takes into account the strength, or morphine equivalent dose (MED), of the prescriptions. Several studies have found that patients with prescriptions that translate to a high average daily morphine equivalent dosage, or MED, (usually above 100-120 mg) are at higher risk for diagnoses of opioid dependence (Sullivan et al. 2010; Edlund et al. 2014) and overdoses (Bohnert et al. 2011). Other indicators for hazardous prescribing focus not only on the quantity of opioids prescribed, but also on a patient's method of obtaining the drugs. "Doctor shopping" and "pharmacy shopping," phenomena in which patients receive opioid prescriptions from multiple prescribers or pharmacies, also correlate with diagnoses of opioid dependence (White et al. 2009), hospitalizations (Jena et al. 2014), and overdose deaths (Hall et al. 2008; Yang et al. 2015).

We therefore construct several measures of potentially risky prescription opioid use. All the measures are at the enrollee-year level and are constructed using the Medicare Prescription Drug Events and the Pharmacy Characteristics files. The level of observation in the Prescription Drug Event file is an "event," or prescription fill, which we map to measures at the enrollee-year level. Each event is associated with an enrollee, the date filled, a national drug code (NDC), a prescribing physician, and the days of supply. We restrict to fills of drugs that contain at least one opioid active ingredient described in the MED conversion table (Appendix Table A.7).

We define and separately analyze three indicator variables as proxies of risky use: High MED, Many Prescribers, and Overlapping Prescriptions. High MED and Many Prescribers are defined following Meara et al. (2016); Overlapping Prescriptions is defined following Logan et al. (2013) and Jena et al. (2014).

Our primary measure of risky use is High MED. This is an indicator for whether average daily MED in any calendar quarter of the year is greater than 120 mg. To define it, we compute the MED for each calendar quarter by multiplying the number of pills by their strength and the morphine equivalence, adding across all fills and ingredients, and then dividing by the number of days in the calendar quarter. In Appendix C, we also explore alternative thresholds for the indicator, including 20 mg intervals between 60 mg and 180 mg. "Many Prescribers" is an indicator for whether an enrollee filled prescriptions associated with four or more unique physicians during the calendar year. "Overlapping Prescriptions" is an indicator for whether the enrollee filled a new opioid prescription before her previous opioid prescription was exhausted. To more effectively target hazardous overlaps, we use an approach similar to prior studies and define this indicator so that it takes the value one only if the second opioid refill was either prescribed by a different doctor (indicating potential doctor shopping) or overlapped with the existing opioid prescription for more than one week (indicating potential use for non-medical purposes). The existing literature has interpreted High MED as a measure of prescription supply, and has interpreted Many Prescribers and Overlapping Prescriptions as measures of doctor shopping.

We also define two other indicators of prescription opioid use for an enrollee-year: Any Opioid Use and Chronic Opioid Use. Following Morden et al. (2014), we define the latter as an indicator for whether the enrollee filled more than six prescriptions in one year.

Finally, we define a number of adverse opioid outcomes. To do so, we must limit our analysis to the 75 percent of non-mover enrollee-years who are not enrolled in Medicare Advantage during the year, so that we can observe their full set of inpatient and outpatient claims. We define a "poisoning event" for each enrollee-year as an ER visit or inpatient hospital admission for poisoning. We define opioid poisonings as

the subset of poisoning events that are so-labelled. We define opioid use disorder as the set of so-labelled diagnoses in the claims data.

Comparison across measures

Among non-movers, about four percent of our enrollee-years satisfy the threshold for High MED, six percent of enrollee-years have Many Prescribers, and about 15 percent of enrollee-years have Overlapping Prescriptions. The pairwise rank correlations among these three proxies of risky opioid use are high (Appendix Table A.8) and all are positively correlated with our other indicators of prescription opioid use.

To evaluate the three measures, we limit our analysis to the non-mover sample with traditional Medicare for which we can fully observe opioid poisonings and opioid use disorders. We further restrict to non-mover enrollee years whom we observe the following year, so that we can examine the relationship between our measures of opioid use in year t and adverse outcomes in year $t+1$ (Appendix Table A.1).

Row (6) shows that, on average, 0.5 percent of our sample is diagnosed with an opioid use disorder each year, and 0.1 percent of our sample is diagnosed with an opioid poisoning event. These outcomes, however, are significantly more prevalent among the beneficiaries that exhibit our measures for risky opioid use in the year preceding. For example, row (5) shows that among the three percent of enrollee-years with High MED, the share with opioid poisonings or opioid use disorders in the next year is about seven and one-half times higher than the baseline. There is substantial variation among these measures; row (6) suggests that these same adverse opioid-related outcomes are only approximately twice as common in the subsample with any opioid prescriptions as compared to the baseline sample. We chose High MED as our baseline measure of risky opioid use because adverse opioid-related outcomes are the most prevalent across all five of the opioid use measures we examined.

Finally, for our preferred measure of risky opioid use—High MED—we examine how geographic patterns of risky use in our population correlate with geographic patterns in adverse opioid outcomes in the overall US population. To do so, we use 2006-2019 state-level data from the CDC Multiple Cause of Death (MCD) File[34] to construct a measure of opioid-related overdose deaths—i.e., deaths due to drug poisoning from heroin, other opioids, methadone, or other synthetic narcotics (such as fentanyl). We define the opioid death rate as these opioid-related deaths as a share of the population. We also measure self-reported opioid misuse rates in the 2006 - 2019 National Survey on Drug Use and Health (NSDUH).[35] At the state level, we calculate that both measures have a high rank correlation with our risky opioid use measure; the average yearly rank correlation is 0.40 for the opioid-related death rate, and 0.40 for the self-reported opioid misuse rate.

[34] Available at https://wonder.cdc.gov/mcd-icd10.html.

[35] This data is available at http://datafiles.samhsa.gov/study-series/national-survey-drug-use-and-health-nsduh-nid13517. From 2006-2014, the survey asked a question about "non-medical use of pain relievers in the past year." This question was removed from the survey in 2015 and replaced with a similar question from 2016-2019 on self-reported prescription pain reliever misuse. We use the share of the adult population who answers "yes" to these questions as our measure of self-reported opioid misuse.

Appendix B: Opioid-Related Policies and Reduced Form Impacts

Although there is growing evidence on a diverse set of opioid-related policies—including Naloxone Access laws and Good Samaritan Laws (Rees et al. 2019), syringe exchange programs (Packham 2022), policies that increase access to treatment for opioid use disorder (Barrette, Dafny, and Shen 2023), and triplicate prescription programs (Alpert et al. 2022) among others—we focused on the three policies that were widely adopted during our sample period and that have been highlighted in the prior literature: Prescription drug monitoring programs ("PDMPs"), pill mill laws, and prescribing limits. We take advantage of the variation across states in whether and when they enacted each policy to estimate their impacts in our study population based on state level difference-in-difference estimates. In this section we describe each policy, our estimates of their impact, and how our results compare to prior analyses of these policies in other populations. Based on these findings, we focus our model-based policy analysis on pill mill laws.

Descriptions of Each Policy

Prescription drug monitoring programs ("PDMPs") are state-level databases that track prescriptions of controlled substances. These provide information on patient prescription histories for physicians as well as providing potentially valuable information for public health officials and law enforcement, all of which is designed to reduce rates of prescribing misuse for monitored substances.[36] Although California established the first PDMP in 1939, the adoption and modernization of PDMPs has been a more recent phenomenon across most states. Modern PDMPs vary widely across states, from simply having established prescription registries, to digitizing them, to mandating their use by physicians. Much of the literature studying these policies focuses on the implementation of "must-access" mandates (Buchmueller and Carey 2018; Alpert, Dykstra, and Jacobson 2024a); these refer to a strengthened set of requirements which dictate that providers access the PDMP prior to prescribing or filling a prescription for a controlled substance. We follow this approach. Our analysis of PDMPs is therefore an analysis of must-access PDMPs.

Pain clinic regulations ("pill mill laws") are laws intended to reduce inappropriate prescribing by a certain set of healthcare facilities referred to as "pain management clinics." These laws typically define these as facilities which engage in the treatment of pain through controlled substances. However, each state's implementation varies slightly in their specific definition of such a facility, and the definition can apply quite broadly to healthcare facilities which aim to treat chronic pain. For example, Texas defines a pain management clinic through a specific set of drugs—opioids, benzodiazepines, barbiturates, or carisoprodol. Kentucky's definition includes all facilities that advertise for any type of pain management services.[37] These laws further define a set of requirements and responsibilities for such facilities. These responsibilities can include explicit certification requirements, requiring physician ownership, complying with specific prescribing restrictions, requiring physician examinations for patients, requiring drug-testing for high-risk patients, and keeping patient records.

Finally, initial opioid prescribing limits ("prescribing limits") represent a recent set of state-level policies targeted towards opioid-naive users with the intention of preventing initiation into addiction and to limit the supply of opioids available for potential diversion (Sacks et al. 2021). These laws typically limit the strength and length of initial prescriptions, with the definition of opioid-naive, maximum prescription thresholds, and maximum prescription durations varying by state.

[36] In particular, the idea is that physicians will adjust downward their prescribing behavior when they identify problematic drug interactions or patients with potentially inappropriate drug-seeking behaviors. However, Baehren et al. (2010) also find evidence that the information can lead to upward adjustments in opioids prescribed when the physician is able to confirm the patient did not have a recent history of controlled substance use. More recently, Alpert, Dykstra, and Jacobson (2024a) consider the hassle costs of required PDMP checks.

[37] These are defined in the Texas Occupations Code § 168.001 and the Kentucky Revised Statutes § 218A.175.

Defining Policy Implementation Dates

We collected data on state-level opioid-related policies from several sources, including legal databases, the Prescription Drug Abuse Policy System (PDAPS), and literature analyzing the effects of opioid-related policies that were enacted during our study period (Buchmueller and Carey 2018; Sacks et al. 2021; Kaestner and Ziedan 2023). We focus on states which implemented these policies during our study period between 2006 and 2019.

One of the challenges for measuring the timing of implementation is defining implementation itself. For example, the PDAPS data reports three sets of dates for different stages of implementation of PDMPs: (a) the passage of the law, (b) when the PDMPs become operational and begin storing data, and (c) when authorized users can actually access data. Despite the variety of potential dates that could be used, the literature studying the effects of these policies presents a mostly-consistent set of dates. When possible, we attempt to follow the literature. After cross-checking implementation dates across sources, we follow the timing used by Sacks et al. (2021) for prescribing limits and must-access PDMPs, and the dates listed on PDAPS for pill mill laws.[38] Appendix Table A.2 presents the resulting implementation dates we use for each policy in each state in our analysis.

Reduced form impact of each policy in our study population

Using our study population of non-movers and the implementation dates described above, we separately estimate the effects of each of the three policies by leveraging their staggered implementation across states. In order to improve power, we allow for year effects to vary flexibly by demographic groups. Specifically, for each policy, we estimate a Poisson event study regression through pseudo-maximum likelihood (Correia, Guimarães, and Zylkin 2020) using the following specification

$$y_{jgt} = \exp\left(\alpha_{jg}^{policy} + \tau_{gt}^{policy} + \sum_{r=-5}^{5} \rho_{r(j,t)}^{policy} \cdot \mathbb{I}_{\{policy_j\}}\right) n_{jgt}\varepsilon_{gjt}, \tag{6}$$

where g denotes a demographic cell defined by the interactions between five-year age bins, race (White, Black, or other), and gender (with demographics defined as in the main text), j denotes states, t denotes calendar years, y_{gjt} denotes the average risky use rate among the state-year-demographic cell, n_{jgt} denotes the number of individuals observed in the state-year-demographic cell, $r(j,t)$ denotes the year relative to the implementation of the relevant opioid policy, and $\mathbb{I}_{\{policy_j\}}$ denotes an indicator variable for whether the state ever passed the relevant opioid policy. The α_{jg}^{policy} fixed effects allow for state fixed effects that vary flexibly by demographic cells, while the τ_{gt}^{policy} fixed effects allow for year effects that likewise vary flexibly by demographic cells. We estimate confidence intervals through 50 iterations of the Bayesian bootstrap procedure, which smooths bootstrap samples through reweighting rather than resampling observations. In particular, we draw fifty sets of weights from a gamma distribution for each state-year-demographic cell, and we re-scale the weight given to each observation by these draws. We then re-estimate $\rho_{r(j,t)}^{policy}$ for each set of drawn weights and compute the standard error across draws.

Appendix Figure A.3 presents our separate estimates of ρ_r^{policy} from equation (6) for must-access PDMPs, pill mill laws, and prescribing limits. The significant upward trend in Panel A makes the effect of must-access PDMPs difficult to interpret, but there is limited evidence of any decrease in risky use. Similarly, Panel C suggests that there were no immediate effects of the policy, and the eventual post-policy decreases

[38] One exception is that we follow Lyapustina et al. (2016) in defining the year of the implementation of Texas's pill mill regulation as 2010 rather than 2009. Although Texas Administrative Code § 192.1, 192.4-192.7 were passed in November 2009 to define pill mills and a set of standards, the bulk of the regulation came shortly afterwards in 2010 through Texas Administrative Code § 195.4 and § 168.001, which required certification, quality assurance procedures, and requirements on ownership.

are noisy. On the other hand, Panel B suggests that pill mill laws had a substantial impact on curbing risky use in our sample, with a gradual decline in risky use beginning in the year after pill mill implementation and growing in the four years following. These findings (which as discussed below, are broadly consistent with the existing literature), motivate us to focus our modeling of policy impacts on the effects of pill mill laws.

The fact that the impact of pill mill laws appears to grow over time could reflect a growing impact of a fully phased-in law, or a gradual phase-in of a law with a time-constant impact, or some combination of both. In practice, pill mill laws are typically not fully implemented immediately after they are enacted: there are often logistical hurdles to implementation as well as subsequent modifications to the policy. Stone et al. (2020) provide a qualitative discussion of factors behind these delays based on interviews with policy-makers in several states that introduced pill mill laws.[39] Appendix Figure A.5 documents the process of gradual modifications by presenting the timing of the initial implementation of pill mill laws and of their subsequent amendments, which we group into seven major categories: (i) explicit penalties for non-compliance, (ii) explicit inspection of pain clinics to ensure compliance, (iii) restrictions on pain clinic physicians (e.g., no prior felonies, no restrictions on their professional licenses, required training), (iv) certification requirements, (v) restrictions on which types of medical professionals are allowed to prescribe opioids, (vi) requirements for drug testing, and (vii) requirements for prescribers to check PDMPs. We document 23 policy amendments for 10 of the 11 states which passed a pill mill law. These frequent policy amendments combined with the qualitative evidence of implementation delays motivate our decision to allow for gradual phase-in of these laws in our model parameterization, as discussed in Section 5.2.

Comparison to existing estimates

Appendix Figure A.4 summarizes our estimates of the impact of each of the three laws on risky use in our sample and compares it with state-level difference-in-differences estimates from four prior studies. For these differences-in-differences estimates, we modify equation (6) to

$$y_{jgt} = \exp\left(\alpha_{jg}^{policy} + \tau_{gt}^{policy} + \text{Post}_{jt} \cdot \mathbb{I}_{\{policy_j\}}\right) n_{jgt} \varepsilon_{gjt}, \tag{7}$$

where Post_{jt} is an indicator for whether $r(j,t) > 0$. Similarly to our policy event study specifications, we estimate confidence intervals for this Poisson regression through 50 iterations of the Bayesian bootstrap procedure. We find that pill mills are associated with a statistically significant 10% decline in risky use, while we estimate substantially smaller and statistically insignificant effects of must-access PDMPs and prescribing limit policies. These findings are broadly consistent with existing estimates where comparisons are feasible.

We identified four other studies that examined the effects of these policies, three of which report results using the same measure of risky use as our baseline measure (average daily morphine-equivalent dose (MED) of more than 120 mg in any calendar quarter). For all of these studies, we report their estimated effect size as a percentage of the mean to make them more easily comparable.

Meara et al. (2016) use individual-level logistic regressions with state and year fixed effects to study the effects of the implementation of opioid-related policies from 2006 through 2012 on various measures of risky use. They study a national sample of beneficiaries enrolled in fee-for-service Medicare Parts A, B, and D, excluding patients with cancer diagnoses, end-stage renal disease, and those in hospice care. Controlling for patient characteristics, state fixed effects, year fixed effects, indicators for the number of laws in each

[39] Potential sources of implementation delays provided by interviewees in Stone et al. (2020) include the time associated with mobilizing pill mill inspections, coordination with law enforcement, information delays associated with definitions and informing providers who "did not understand that they needed to register as a pain clinic to comply with the law", and the time required for external parties to raise flags about signals of potential non-compliance with the law.

state, and the further indicators for how many laws were added since 2006, they find limited systematic evidence of strong effects for any of the policies that we study on any measure of risky use, although their point estimate on the effect of pill mills is large and negative (but noisy). Appendix Figure A.4 presents their results from the point estimates and standard errors presented in Table S9 on the marginal effects of each policy on the same measure of risky use that we use.

Buchmueller and Carey (2018) use a differences-in-differences design to study the effect of PDMPs implementation between 2007 and 2013 on measures of risky use among a 5 percent subsample of Medicare beneficiaries enrolled in Part D and fee-for-service Medicare. They carefully distinguish between PDMPs with must-access provisions and those without any. Although they find statistically significant declines in risky use using doctor-shopping related measures, such as probability of beneficiaries obtaining opioids from five or more prescribers or five or more pharmacies, they estimate a fairly precise zero on the effect of these policies on our measure of risky use. Appendix Figure A.4 presents these results from their baseline specification in Table 6, Panel B.

Sacks et al. (2021) also uses a differences-in-differences design that allows for both must-access PDMPs and prescribing limits to have constant and additively separable effects on risky use outcomes within a commercially insured population between 2007 and 2018. They find that while reducing the size of the initial prescriptions, prescribing limits may have actually increased the number of new opioid prescriptions on the extensive margin. However, they find little effect for either policy on a wide variety of potential measures of risky use: high dosages, overlapping claims, or doctor shopping outcomes. Appendix Figure A.4 presents their results from Table 5, Panel A, on the effects of both policies on our measure of risky use.

Finally, Kaestner and Ziedan (2023) study the effects of "modern" PDMPs and pill mills using aggregate data on prescription opioid sales from the DEA's Automation of Reports and Consolidated Orders System (ARCOS).[40] They do not construct measures of risky use, so we use their results on aggregate sales—measured in per-capita morphine equivalents grams (MEG)—as our benchmark. This measure accounts for heterogeneity in the strength of different drugs by weight. They use state-level event studies to measure the effects of both policies on aggregate opioid sales per capita. Although this study does not use individual-level data on prescription opioid receipt, the more recent time window and the comparison of effect sizes across policies was helpful in choosing a policy to focus on in our paper. Appendix Figure A.4 presents their results on these effects from the second and fourth column in their Table 2.

Appendix Figure A.4 shows that in most cases, the studies report wide confidence intervals, but that the strongest evidence of an effect is for pill mill laws. In particular, Kaestner and Ziedan (2023) estimate that pill mill laws reduced sales of prescription opioids per capita by a statistically significant 12 percent while finding a smaller and statistically insignificant effect for must-access PDMPs. In addition, while Meara et al. (2016) estimate statistically insignificant impacts for all three policies, the point estimate is largest for pill mill policies. The other two studies do not look at pill mill laws, and find statistically insignificant impacts for the other two policies, with point estimates smaller than the pill mill estimates in Kaestner and Ziedan (2023) and Meara et al. (2016).

[40] They define "modern" PDMP policies as those that implement a directory that is accessible on demand by authorized users. These are a superset of MA-PDMPs, and more information on their definition can be found in their Appendix Table 1.

Appendix C: Additional Analyses

Comparison of patterns in our sample to overall U.S. population

We compare national trends and geographic patterns of opioid prescription rates in our population to the general US population. To do so, we obtain data on opioid prescription fills per capita from county- and state-level averages of QuintilesIMS opioid prescription data, which are made publicly available by the Centers for Disease Control and Prevention. QuintilesIMS collects data on prescriptions based on a sample of 59,000 retail pharmacies, which collectively dispense nearly 88 percent of all prescriptions in the US (CDC 2017). The aggregated QuintilesIMS data set contains the number of opioid prescriptions per capita in each year.[41] We define an identical variable in our data, the number of opioid prescription fills per capita (Opioid Fills). Both are defined over the 2006 - 2019 time period.

Appendix Figure A.14 shows that national trends in opioid prescriptions per capita have evolved similarly in our sample and in the general US population, although with a substantially higher level for the disabled population. Our measure of opioid fills for our SSDI population and the QuintilesIMS measure of opioid fills for the general population are also highly correlated across geographies, with a correlation coefficient of 0.80 at the state level.

Robustness of Reduced Form Evidence from Moves

The results in Section 4 on risky opioid use are robust to several alternative specifications, including defining the geographic areas of interest at different levels, excluding individual and relative year fixed effects, and defining our sample in alternative ways to handle entry or exit by enrollees during our sample period.

Appendix Figure A.15 presents our event study results when estimated at various geographic levels (coefficients in Appendix Table A.9). Panel A shows our baseline results, with state as the geographic unit, for ease of comparison. We find that our event study results are similar if we instead define the geographic areas of interest—and hence redefine who is a mover—at the commuting zone level (Panel B) or the county level (Panel C), rather than the state level. Commuting zones are collections of counties defined to approximate local labor markets; there are about 700 commuting zones and 3,000 counties in the United States.

We also consider the robustness of our analysis to alternative ways the literature has measured risky opioid use. As we discussed in Appendix A, two alternative measures of risky use in the literature are the Many Prescribers and Overlapping Prescriptions measures. We show here that the basic event study results from estimating equation 3 are similar if we use these alternative measures instead of our baseline measure, High MED. Appendix Table A.10 presents summary statistics on these measures, Appendix Table A.11 summarizes the results, and Appendix Figure A.16 shows the corresponding event studies. The immediate discrete jump upon move using Overlapping Prescriptions is similar to that of our baseline measure, while the jump for Many Prescribers is larger. The post-move convergence is somewhat less pronounced with these other measures.[42] We also consider alternative ways of defining High MED by exploring other average daily MED thresholds. Appendix Figure A.17 shows that our estimated coefficients are extremely stable across these alternative thresholds, and that the event study dynamics are quite similar when we instead use average daily MED as a continuous outcome.

[41] Available at https://www.cdc.gov/drugoverdose/maps/rxrate-maps.html.

[42] A lack of post-trends is plausible within our model framework. In Section 4.2, we analyzed a case (Case 4) in which place-based addiction transitions were higher in high risky use locations. However, if we relax this assumption—consider, for example, if high risky use rates were primarily driven by person-specific factors or availability rather than place-based addiction factors—the average risky use rates in a cohort could converge in the opposite direction of δ_{cr}. Thus, we could see zero or even negative post-trends.

Appendix Figures A.7 and A.18 show robustness to modifications in our event study specification, summarized in Table A.12. In Appendix Figure A.7, we define the dependent variable in levels, rather than in differences; in other words, instead of defining $\hat{T}_{cr}^{X}$ to be the average of the difference in differences $(y_{ir} - y_{i,-1}) - \left(\hat{m}_{cr}^{X} - \hat{m}_{c,-1}^{X}\right)$ among the set $\mathscr{I}_{c}^{X}$ of movers in cohort c with characteristics X, we define it as $\left(y_{ir} - \hat{m}_{cr}^{X}\right)$. As in our baseline specification, we then let $\hat{T}_{cr}$ denote the average of $\hat{T}_{cr}^{X}$ across X, weighting by the number of movers in c with each set of characteristics X. With this alternative specification, we no longer adjust for any fixed individual differences in propensity for risky use. The figure shows that we continue to find little systematic pre-move trend, and a similar post-move pattern of coefficients. However, the level of the pre-move coefficients is now about 0.15, indicating that there is potential selection in where movers move. In particular, the results indicate that movers to destinations with a higher propensity for risky use than in their origin tend, prior to their move, to have a higher propensity for risky use than matched non-movers in their origin.

Appendix Figure A.18, Panel A presents our results if we exclude relative year fixed effects from estimation of equation (3). Our baseline specification constructs the dependent variable $\hat{T}_{cr}$ in equation (3) by averaging across $\hat{T}_{cr}^{X}$ across movers in cohort c, where X consists of five-year age bins (as measured by birth year), race, and gender. Panel B presents our results when we omit any conditioning on covariates X in constructing $\hat{T}_{cr}$. We continue to observe the key post-move features of our event study: a discrete jump upon move and then post-trends.

Next, we consider the robustness of our results to alternative ways of handling the entry and exit of individuals from our sample. Individuals may not be observed for all years because of entry and exit on two margins. The first margin is SSDI: individuals become eligible and enroll for SSDI in different years, and exit either due to losing eligibility or death. The second margin is Part D coverage, and individuals may enter and exit our sample by enrolling or dis-enrolling in Part D coverage for each full calendar year. We therefore estimate event studies where we adjust our sample along both of these margins.

First, Appendix Figure A.19 suggests that our event study results hold in a balanced panel of individuals who are observed in our sample for three years pre- and post-move, indicating that these features are not simply driven by the changing composition of groups at each length of exposure to the new location. Next, we turn to specifications which adjust our sample on non-SSDI margins. We present these results in Appendix Figure A.20, and summarize them in Appendix Table A.14. Our baseline sample restricts to enrollee-years where enrollees are covered for the full year (i.e., have 12 months of Part D coverage). A given enrollee may therefore enter or exit the sample if she gains or loses just a few months of Part D coverage. Panel A therefore shows results if we further restrict the sample to enrollees who had Part D coverage for the full year (all 12 months) in all years they appear in our sample, while Panel B restricts to enrollees who are alive in all years during our study period.

We also show that our results are not driven by only the most common moves or by patients who might be expected to have very high opioid use. In Figure A.20, Panel C, we exclude moves to or from the three most common destinations in our sample (Florida, Texas, and California), and in Panel D, we exclude from the sample all enrollees for whom we observe a cancer diagnosis or hospice stay. Appendix Figure A.6 also shows that our results are robust to restricting to moves to either higher risky use locations or lower risky use locations, and these results are summarized in Appendix Table A.13.

Finally, to further explore the potential impacts of selection into our sample, Appendix Figure A.21 estimates equation (3) using as the dependent variable whether the enrollee-year had 12 months of Part D coverage. For all of these analyses, $\hat{\delta}_{cr}$ is still defined based on risky opioid use. The results show some evidence of selection into Part D upon move; for example, a move from a 20th percentile risky opioid use state to a 80th percentile state is associated with about a two percentage point increase in the probability of having Part D coverage (relative to a mean of 80 percent). Reassuringly, however, Appendix Table A.14 shows that restricting attention to individuals whose selection along this margin does not change during the

sample does not affect our main results.

Commuting Zone-Level Model Estimation

We re-estimate our model at the commuting zone level in order to leverage additional variation in geographic correlates. For computational feasibility, however, we begin by restricting to the 100 largest commuting zones in our sample[43]. We further restrict to commuting zones with at least 50,000 movers—including movers coming both from and to the commuting zone—and we remove the few remaining commuting zones that span multiple states and would have been "partially-treated" by state policies. This leaves us with a sample of 91 of the largest remaining commuting zones, spanning 36 states. We do not re-estimate our policy parameters on this sample because the treated population now represents a subsample of both our treated states and a subsample of geographies within our treated states. Furthermore, even with these restrictions, we now estimate 475 parameters — compared to 282 parameters in our baseline model.

Appendix Figure A.22 provides a map of the included commuting zones. We provide summary statistics on this estimation sample below in Appendix Table A.15. Comparing these moments to the moments constructed from our baseline sample, there are nearly 50% more moments — driven by the larger number of geographic units and pairs — but these moments are constructed using a sample that is approximately 50% smaller. Nonetheless, we show in Figure A.23 that our model predictions align well with the mover and non-mover moments. Finally, row (2) of Appendix Table A.5 suggests that the relative importance of each set of parameters is similar to our baseline model.

Medicare Elderly Sample: Reduced form evidence from moves and model estimates

We replicate our baseline set of reduced form and model estimation results for a sample of elderly Medicare patients. In order to construct this sample, we begin with the same 20 percent sample of all Medicare enrollees. Rather than restricting our sample to individuals enrolled in SSDI, however, we now instead limit our sample to patient-years in which the patient was at least 65 years old. We then simply replicate the remaining steps in Section 2.2 and define a set of movers and non-movers for this sample.

Appendix Table A.16 reports summary statistics for the Medicare elderly sample. Relative to our baseline sample in Table 1, both movers and non-movers are more likely to be white and female. They are significantly less likely to receive Medicaid, and by construction, the average age in our sample is much older. Most strikingly, the rates of opioid use in this sample are lower. While about one-fifth of individuals receive a prescription for any opioids in a year, only about four-tenths of a percent exhibit risky use—approximately ten times fewer than in the baseline sample.

Despite the differences in these samples, when we apply our empirical strategy outlined in 4.1, we find strikingly similar results as our baseline sample. Appendix Figure A.24 shows estimates of $\hat{\mu}_r$ as estimated on the Medicare elderly sample. We find strikingly similar results to our baseline specification. We find little systematic trend in the coefficient before the move but observe an immediate jump in risky use upon move and a gradual post-move convergence. Furthermore, as column (1) of Appendix Table A.17 suggests, even the magnitudes of the jump and the post-move convergence are quite similar to our baseline specification.

Appendix Figure A.25 shows these results as estimated separately for movers with opioid utilization in the year prior to move ("prior users") and movers without opioid use in the year prior to move ("opioid naives"). These also show the same qualitative patterns as our baseline sample. There is an immediate jump upon moves for prior users, while there is little change upon move for opioid naives. In the subsequent post-move years, risky use rates for prior users and opioid naives both increase gradually, although the latter

[43] If we were to estimate our model with all 709 commuting zones, the estimation sample would only be 25 percent larger but have approximately 1,300 percent more parameters.

post-trend is smaller and noisy. Columns (2) and (3) of Appendix Table A.17 suggest that relative to the baseline, the magnitudes of the effects are smaller for opioid-naives and larger for prior users.

Next, we re-estimate our model using moments derived from the elderly Medicare sample. Appendix Table A.18 provides an overview of these moments. Appendix Figure A.26 shows the geographic distribution of factors affecting availability, place-based addiction transitions, and person-based addiction transitions. The geographic distribution of these parameters looks quite similar to our baseline sample. We also find in Appendix Figure A.27 that our model replicates non-mover moments from our estimation sample very well, although the fit on mover moments is a bit worse.

Next, we use these model estimates to replicate the mover event study design, and Figure A.28 presents our results. We find once again that the model delivers predictions that are largely consistent with the mover variation that we estimate, although the standard errors are quite large due to the low frequency of risky use in this population. Finally, the correlations from Appendix Figure A.26 suggest that the role of geographic variation in place-based availability and place-based addiction transitions seems to be larger than the role of geographic variation in person-based addiction transitions. When we replicate the exercise detailed in Section 6.2 to examine the importance of these channels, we find results consistent with this; Appendix Table A.5, row (3) suggests that place-based factors play a larger relative role compared to person-based factors in the Medicare elderly sample compared to our baseline SSDI sample.

Appendix D: Proof of Proposition 1

We begin by showing that under Assumption 1, $\mathbb{E}_{i\in\mathscr{I}_c}\left[\hat{T}_{cr}\right] = T_{cr}$. To see this, we note that the average treatment effect among movers with characteristics X is:

$$\begin{aligned}
T_{cr}^X =& \mathbb{E}_{i\in\mathscr{I}_c^X}\left[y_{ir}\left(\mathbf{h}_{cr}\right) - y_{ir}\left(\mathbf{h}_{cr}^0\right)\right] \\
=& \mathbb{E}_{i\in\mathscr{I}_c^X}\left[\left(y_{ir}\left(\mathbf{h}_{cr}\right) - y_{i,-1}\left(\mathbf{h}_{c,-1}^0\right)\right) - \left(y_{ir}\left(\mathbf{h}_{cr}^0\right) - y_{i,-1}\left(\mathbf{h}_{c,-1}^0\right)\right)\right] \\
=& \mathbb{E}_{i\in\mathscr{I}_c^X}\left[y_{ir}\left(\mathbf{h}_{cr}\right) - y_{i,-1}\left(\mathbf{h}_{c,-1}^0\right) - \left(\hat{m}_{cr}^X - \hat{m}_{c,-1}^X\right)\right] \\
=& \mathbb{E}_{i\in\mathscr{I}_c^X}\left[y_{ir} - y_{i,-1} - \left(\hat{m}_{cr}^X - \hat{m}_{c,-1}^X\right)\right] \\
=& \mathbb{E}_{i\in\mathscr{I}_c^X}\left[\hat{T}_{cr}^X\right]
\end{aligned}$$

where the third line follows from Assumption 1 and the fourth line follows from the facts that $\mathbf{h}_{c,-1}^0 = \mathbf{h}_{c,-1}$ and the fact that $y_{ir}\left(\mathbf{h}_{it}\right)$ is equal to the actual observed outcomes y_{ir} for all r. Averaging over X then implies $\mathbb{E}_{i\in\mathscr{I}_c}\left[\hat{T}_{cr}\right] = T_{cr}$. Next, to show that $\hat{T}_{cr}$ is also a consistent estimator for T_{cr}, we recall that $\hat{T}_{cr}^X$ is simply the sample average of $y_{ir} - y_{i,-1} - \left(\hat{m}_{cr}^X - \hat{m}_{c,-1}^X\right)$ among movers in cohort c with characteristics X. It follows that $\hat{T}_{cr}^X \to_p T_{cr}^X$, and averaging once again over X implies that $\hat{T}_{cr} \to_p T_{cr}$.

Next, we would like to show that the event study coefficients $\hat{\mu}_r$ from equation (3) converge in probability to

$$\mu_r = \sum_c w_c \frac{T_{cr}}{\delta_{cr}}.$$

Observing that the coefficient for each relative year may simply be written as the result of a regression through the origin, we expand the coefficient as follows

$$\hat{\mu}_r = \frac{\sum_c |\mathscr{I}_c| \hat{\delta}_{cr}\hat{T}_{cr}}{\sum_c |\mathscr{I}_c| \hat{\delta}_{cr}^2} \to_p \frac{\sum_c |\mathscr{I}_c| \delta_{cr}T_{cr}}{\sum_c |\mathscr{I}_c| \delta_{cr}^2},$$

where the convergence in probability follows from the convergence of the sample analog $\hat{\delta}_{cr} \to_p \delta_{cr}$ and treatment effect $\hat{T}_{cr} \to_p T_{cr} + \rho_r$, as shown above.

Appendix E: Estimation Details

Generalized Method of Moments

The model, as parameterized in Section 5, includes five sets of parameters which vary by location—γ_j, π_j^+, π_j^-, η_j^+, η_j^-—as well as the level of selection s, the separable temporal component of availability τ_t for each year of our sample, the impacts of moving upon availability ρ_r for the five years before and five years after the move, the policy parameters ζ_{jt}^{π} and ζ_{jt}^{γ}, and the five policy implementation parameters $share_t$.

In order to avoid collinearity in the sums that are used in the computation of addiction ($\pi_j + \eta_j$) and availability ($\gamma_j + \tau_t + \rho_r$) in our model, we normalize three sets of values to zero as references: the first calendar year availability parameter, the effect of the fifth-year pre-move on availability, and an arbitrary state's person-specific addiction transition effects. The vector of parameters to be estimated therefore consists of 287 values and is given by $\theta = \left\{\gamma_j, \pi_j^+, \pi_j^-, \eta_j^+, \eta_j^-, s, \tau_t, \rho_r, \zeta_{jt}^{\pi}, \zeta_{jt}^{\gamma}, share_t\right\}$.

We construct our score function as follows. We first construct two vector-valued functions for average risky use rates as a function of our model parameters. For movers, we construct the function $M^{mover}(\theta)$ where cr indexes each observed combination of cohort and relative year and each entry $M_{cr}^{mover}(\theta)$ is given by the average cohort risky use rate $\bar{y}_{cr}$ as calculated according to Section 5. For non-movers, we construct the function $M^{non-mover}(\theta)$ where jt indexes each state-year, and each entry $M_{jt}^{nonmover}(\theta)$ is given by the average state-year risky use rate as calculated according to Section 5. Next, we allow the vector $Z_i^{mover} := \{(Y_{ir}, X_{ir}')'\}_r$ to be the vector collecting the variables for individuals i across relative years r, where

$$Y_{ir} = \begin{cases} y_{ir} & \text{if } i \text{ is a mover observed in relative year } r \\ 0 & \text{else} \end{cases}$$

denotes the risky use outcome of the individual in any given relative year and X_{ir} is a vector with a 1 in the index cr where the individual-year combination ir is a subset of cohort-year ct and a 0 in all other entries. We define the vector $Z_i^{nonmover}$ analogously for non-movers, such that $Z_i^{mover} := \{(Y_{it}, X_{it}')'\}_t$ collects variables for individuals i across calendar years t. Similarly, Y_{it} denotes if they are non-movers observed in calendar year t, and X_{it} is a vector with a 1 in the index jt when the individual-year combination is a subset of the state-year jt. We stack these vectors to form Z_i. Finally, for each of our observed combinations cr, we construct a corresponding function g_{cr} as

$$g_{cr}(Z_i, \theta) = [(Y_{ir} - M_{cr}^{mover}(\theta))] X_{ir}' W_{cr} + \left[\left(Y_{it} - M_{jt}^{nonmover}(\theta)\right)\right] X_{it}' W_{jt},$$

where W_{cr} is a vector with 1 in the entry corresponding to cr and a 0 everywhere else, and W_{jt} is a vector with 1 in the entry corresponding to jt and 0 everywhere else. We stack these functions g_{cr} into the score function $g(Z, \theta)$, where Z denotes a generic observation from the population.

Using $g(\theta) \equiv \mathbb{E}g(Z, \theta)$ to denote the expectation of this score function with respect to the generic observation from the population, we then define the corresponding empirical estimator $\hat{g}(\theta)$ as $\frac{1}{N}\sum_{i=1}^{N} g(Z_i, \theta)$ where Z_i now denotes an observation from our sample. The GMM estimator is therefore defined by

$$\hat{\theta} = \arg\min_{\theta} \hat{g}(\theta)' \hat{A} \hat{g}(\theta),$$

where we choose the weighting matrix $\hat{A}$ to be the diagonal matrix where each diagonal entry is the number of individuals in our sample who were observed in the relevant cohort-year or state-year, re-weighted in aggregate such that non-movers receive the square root of the total number of non-mover years observed and movers receive the square root of the total number of mover-years observed. We choose this weighting matrix because the typical two-step procedure is ill-suited for the structure of our problem, which includes

moments that correspond to very few observations in our data.

Parameter Space Constraints

We constrain our parameter space in two ways. First, as our parameters describe probabilities and rates, we require that all addiction and availability parameters that an individual could be subject to in the model are bounded between 0 and 1. In practice, we find that this constraint does not bind for availability. The constraint binds at 0 for 0.3% of possible pairwise combinations of states for the probability of transitioning into addiction, and it binds at 0 for 3% of pairwise combinations of states for the probability of transitioning out of addiction. Second, we restrict our parameter space such that the combination of parameters generates a weighted global mean share of addiction equal to the specified value $\bar{a} = 0.10$. Specifically, we require our estimation to satisfy

$$\bar{a} = \sum_{t=2006}^{2019} \sum_{j \in \mathscr{J}} \hat{\bar{a}}_{jt} \frac{\hat{n}_{jt}}{N_{nm}},$$

where $\hat{n}_{cr}$ denotes the number of enrollee-years used to construct the moment $\hat{\bar{a}}_{jt}$, $\bar{a}_{jt}$ denotes the average share of addiction for non-movers in a given state-year (as calculated in Section 5), and N_{nm} is the total number of non-mover enrollee-years in the sample.

We view this constraint as a normalization because the multiplicative property of our outcome implies that the relative magnitudes of addiction and availability parameters are not meaningful or well-identified. For example, doubling the shares of addiction and halving all availability parameters would leave predicted risky use unchanged in our model. While our specific parameterizations mean that it is not always possible to directly change all addiction or all availability parameters in such a way, this restriction is driven by our functional form assumptions. Thus, we instead focus on the meaningful distribution of parameters within the addiction and availability channels, and we simply set a global mean share of addiction.

We therefore draw our global mean share of addiction among all individual-years from estimates of non-medical prescription pain reliever (PPR) use in the Department of Health and Human Services' National Survey on Drug Use and Health (NSDUH). Estimates from the mid-year of our sample (the 2011 NSDUH) suggest the prevalence of non-medical PPR use was roughly 4% over a 12-month period and 13% over the lifetime of the individuals surveyed (estimates are nearly identical within the small set of SSDI recipients surveyed). Noting that our model implies that the share of addicted individuals must be greater than or equal to the observed rate of risky opioid use and that individuals transition in and out of addiction, we choose 10% as a rough anchor for addiction between annual and lifetime non-medical opioid use. We also show the robustness of our parameter estimates to various other values of $\bar{a}$ in the robustness subsection below.

Estimation Implementation

Our actual estimation occurs in three steps. In the first step of our estimation, we set various initial values for the optimization algorithm. Using sixteen sets of initial values, we set parameter values equal to a constant ranging between 0.00 and 0.15, including the end-points and evenly distributed between. Second, we minimize the weighted distance between the empirical and predicted moments by using the Levenberg-Marquardt algorithm separately for each of these initial values (Moré 1978). We use the penalty approach to non-linear optimization in order to enforce our parameter space constraints. The algorithm converges when the consecutive values of the sum of squares in the objective function are less than 10^{-10} apart. Finally, we select the globally optimal solution as our solution.

Monte Carlo Exercises

Although our estimators are consistent, given the large number of parameters that we must estimate for our model, we also carefully examine the finite-sample properties of our estimators by conducting Monte Carlo simulations. Using the results from Section 5.4 as the true set of parameters, we simulate the risky use of each estimation sample fifty times. We first confirm that as we decrease the noise in our simulated moments by increasing the number of beneficiaries in each sample, our estimated parameters approach the true parameters. We also find that although the estimators for each individual parameter are quite noisy in our finite sample, our key counterfactual statistics are quite stable. In particular, we use the Monte Carlo simulations to assess the consistency of our key counterfactual results. We find evidence in Appendix Tables A.19 and A.20 that suggests that the effects of finite-sample bias on our key counterfactuals are limited and unlikely to affect our conclusions.

Appendix F: Underlying Model of Substitution

Intuition for Learning About the Strength of Substitution Towards Illicit Opioids

We investigate the relationship between our estimated channels for risky prescription opioid use—addiction and availability—and adverse opioid-related outcomes. In particular, we consider how this relationship might be affected by substitution towards illicit opioid use. We do not observe illicit opioid use, nor do we attempt to measure a proxy for it. Rather, we consider three hypotheses in a simple modeling space around the relationship between addiction, the availability of risky opioid prescriptions, and illicit opioid use. We discuss the predictions of each of these models for the size and magnitude of correlations between our estimated model parameters and rates of three opioid-related adverse outcomes. While we cannot directly estimate the rest of these hypothesized relationships, we outline three testable hypotheses where we vary: (i) the direct relationship between risky prescription opioid use and illicit opioid use, (ii) the direct relationship between risky prescription use and adverse opioid-related outcomes.

We first visualize our key hypothesized relationships through a directed acyclic graph structure presented in Appendix Figure A.29. Panel A presents an overview of the limited modeling space that we consider. In this graph, we take two sets of relationships as given. First, the availability and addiction parameters that we estimate from our model are mechanically predictive of area-level rates of risky prescription opioid use. Second, we assume that risky prescription opioid use and illicit opioid use are both positively predictive of adverse opioid-related outcomes.

The key relationships of interest are the degrees to addiction and risky prescription opioid use drive substitution towards illicit opioid use. These relationships are shown through the two dashed red lines, and our three hypotheses are edge cases about the nature of these relationships. Each of these three hypotheses represent different degrees of substitution, ranging from zero substitution to substitution along both channels.

Panel B shows our first hypothesis which considers a model where there is no substitution through either channel. Under this hypothesis, neither addiction nor risky use directly affect illicit opioid use. Thus, variation in addiction and availability parameters should only affect adverse outcomes through their effects on risky prescription opioid use. The ratio of the availability and addiction coefficients in a regression on adverse opioid-related outcomes should be *positive and proportional* between the risky use specification and the specifications with other adverse outcomes.

Panel C shows our second hypothesis under which there is substitution through only the addiction channel: addicted individuals use risky prescription opioids and illicit opioids, but changes in risky prescription opioid usage do not directly affect illicit opioids. Under this hypothesis, availability still drives adverse opioid-related outcomes through its effects on risky use, but addiction drives adverse outcomes through both channels. We should therefore estimate availability and addiction coefficients that are *positive but non-proportional.*

Finally, Panel D considers a third hypothesis under which risky prescription opioid use and illicit opioid use are also direct substitutes. In this hypothesis, there is a direct positive relationship between addiction and illicit opioid use, while there is a direct negative relationship between risky prescription opioid use and illicit opioid use. Under this hypothesis we also assume that decreases in risky prescription opioid use are fully offset by increases in illicit opioid use, which we dub "full substitution." In this case, as long as illicit opioid use is more likely to result in adverse outcomes than risky prescription opioid use, we should estimate a *zero or negative coefficient on availability.*

Formalizing the Hypotheses

Next, we formalize these hypotheses by using the following flexible functional forms to describe these relationships:

$$\begin{aligned} y_j &= \gamma_j \bar{a}_j \\ i_j &= \zeta^{illicit}(y_j, \bar{a}_j) + \varepsilon_j^{illicit} \\ d_j &= \nu^{adverse}(y_j, i_j) + \varepsilon_j^{adverse} \end{aligned} \quad , \tag{8}$$

where j denotes geographies, i_j denotes illicit opioid use, d_j denotes adverse outcomes, and $\varepsilon_j^{illicit}$ and $\varepsilon_j^{illicit}$ denote orthogonal and unobserved drivers of illicit opioid use and adverse opioid-related outcomes that are outside of the model. The two unobserved functions of interest—and the subject of our three hypotheses—are $\zeta^{illicit}(\gamma_j, \bar{a}_j)$ and $\nu^{adverse}(y_j, i_j)$. For simplicity, we assume that both functions are monotonic in all inputs. We also assume that illicit opioids are more likely to drive opioid-related adverse outcomes than risky prescription opioid use such that $\frac{\delta \nu^{adverse}}{\delta i_j} > \frac{\delta \nu^{adverse}}{\delta y_j} > 0$.

Using this notation, we can define our three hypothesis formally as follows:

Hypothesis 1. *(No substitution).* Illicit use does not depend on either risky use or addiction, such that

$$\frac{\delta \zeta^{illicit}(y_j, \bar{a}_j)}{\delta y_j} = 0, \frac{\delta \zeta^{illicit}(y_j, \bar{a}_j)}{\delta \bar{a}_j} = 0 \quad \forall \gamma_j, \bar{a}_j.$$

This implies that a relationship between adverse opioid-related outcomes and the estimated availability and addiction parameters that can be expressed as

$$\frac{\frac{\delta d_j(y_j, i_j)}{\delta \gamma_j}}{\frac{\delta d_j(y_j, i_j)}{\delta \bar{a}_j}} = \frac{\frac{\delta y_j}{\delta \gamma_j}}{\frac{\delta y_j}{\delta \bar{a}_j}}. \tag{9}$$

Hypothesis 2. *(Substitution along addiction channel).* Illicit use depends only on addiction such that

$$\frac{\delta \zeta^{illicit}(y_j, \bar{a}_j)}{\delta y_j} = 0, \frac{\delta \zeta^{illicit}(y_j, \bar{a}_j)}{\delta \bar{a}_j} > 0.$$

This implies that a relationship between adverse opioid-related outcomes and the estimated availability and addiction parameters that can be expressed as

$$\frac{\frac{\delta d_j(y_j, i_j)}{\delta \gamma_j}}{\frac{\delta d_j(y_j, i_j)}{\delta \bar{a}_j}} < \frac{\frac{\delta y_j}{\delta \gamma_j}}{\frac{\delta y_j}{\delta \bar{a}_j}}. \tag{10}$$

Hypothesis 3. *(Substitution along both channels).* Illicit use depends both on risky use and addiction, and decreases in risky prescription opioid use are fully offset by increases in illicit opioid use such that

$$\frac{\delta \zeta^{illicit}(y_j, \bar{a}_j)}{\delta y_j} \leq -1, \frac{\delta \zeta^{illicit}(y_j, \bar{a}_j)}{\delta \bar{a}_j} > 0.$$

Combined with our assumption that $\frac{\delta \nu^{adverse}}{\delta i_j} > \frac{\delta \nu^{adverse}}{\delta y_j} > 0$, this implies that opioid-related adverse

outcomes are decreasing in availability:

$$\frac{\delta d_j(y_j, i_j)}{\delta \gamma_j} < 0. \tag{11}$$

Testing the Hypotheses

Finally, we test the hypotheses from equations (9), (10), and (11) by regressing our addiction and availability coefficients on both risky use and opioid-related adverse outcomes. Table 5 presents our estimates.

Our first hypothesis predicts that if there is substitution at all—and neither addiction nor risky use directly affects illicit opioid use—then the estimated coefficients in columns (2) - (4) should have the same ratios as in column (1). We find instead that the ratio of availability coefficients to person-based addiction and place-based addiction coefficients is significantly lower for our adverse outcomes than for risky use. The third hypothesis, on the other hand, predicts that availability will have a negative coefficient. We find that the availability of prescription opioids for risky use still has meaningfully large and positive coefficient. A one standard-deviation increase in availability is associated with a 10% increase in opioid-related poisonings and a 4% increase in any drug-related poisoning. Naturally, this does not preclude versions of this third hypothesis where there is less than full substitution and reductions in risky prescription opioid use are not fully offset by illicit opioid use.[44] Thus, our evidence is most consistent with our second hypothesis that there is less than full substitution: availability is less predictive of opioid-related adverse outcomes than addiction coefficients, but the coefficient is still positive and significant.

[44]This "less than full substitution" would be given by $\frac{\delta \zeta^{illicit}(y_j, \bar{a}_j)}{\delta y_j} \in (-1, 0)$.

Appendix G: Alternative Specifications

Appendix Table A.5 shows the results from repeating the exercise detailed in Section 6.2 across several extensions and modifications of our model. In this section, we provide more detail on each of these extensions.

The Interactions Between Place-Based and Person-Based Addiction Parameters

We allow the model to endogenously estimate how place-based and person-based addiction parameters interact through a CES-style function. We denote the probability of transitioning into addiction by $\left(\left[\pi_{jt}^{+}\right]^{\alpha}+\left[\eta_{i}^{+}\right]^{\alpha}\right)^{\frac{1}{\alpha}}$ and the probability of transitioning out of addiction by $\left(\left[\pi_{jt}^{-}\right]^{\alpha}+\left[\eta_{i}^{-}\right]^{\alpha}\right)^{\frac{1}{\alpha}}$, where the π and η terms represent place- and person-specific factors respectively, and α determines how they interact. Thus, we can rewrite the probability of addiction more generally as:

$$Pr(a_{i,t}=1|j(i,t)=j)=\begin{cases}1-\left(\left[\pi_{jt}^{-}\right]^{\alpha}+\left[\eta_{i}^{-}\right]^{\alpha}\right)^{\frac{1}{\alpha}} & \text{if } a_{i,t-1}=1\\ \left(\left[\pi_{jt}^{+}\right]^{\alpha}+\left[\eta_{i}^{+}\right]^{\alpha}\right)^{\frac{1}{\alpha}} & \text{if } a_{i,t-1}=0.\end{cases}$$

This CES-style function flexibly encompasses many potential interactions between person-based and place-based addiction factors. For example, $\alpha=1$ corresponds to our baseline model where these factors are additively separable, while $\alpha\to 0$ corresponds to a multiplicative form that exhibits supermodularity. Furthermore, when $\alpha\to-\infty$ or $\alpha\to\infty$, the probability of transitioning into and out of addiction become $\min(\pi_{jt},\eta_i)$ and $\max(\pi_{jt},\eta_i)$ respectively. Results from repeating the exercise detailed in Section 6.2 on these specifications are presented in rows (4) - (6) of Appendix Table A.5. When we estimate α under alternative specifications, we restrict the domain of these parameters to [0, 1] such that all addiction transitions can be interpreted as probabilities.

Allowing for Risky Use by Non-Addicted Individuals

We allow for individuals who are not addicted to receive risky prescriptions with a constant probability across time and geographies. Denoting this probability as λ, we have that the average risky use rate for cohort c and relative year r becomes

$$\bar{y}_{cr}=\bar{a}_{cr}\cdot\gamma_{cr}+(1-\bar{a}_{cr})\cdot\lambda.$$

In this extension, we endogenously estimate this parameter λ as an additional parameter. We estimate that the probability of risky use among non-addicted individuals is 1.5%, while the average probability of risky use among addicted individuals in the same model is 21.4%, nearly fifteen times higher. This non-zero probability of risky use among non-addicted individuals means that policies aimed towards reducing place-based availability (γ_j) and place-based addiction (π_j) are mechanically more likely to have smaller effects, as row (7) of Appendix Table A.5 supports. The *relative* importance of each channel however, remains similar to our baseline model.

Addiction Transitions During The Year of Move

Our baseline model assumes that no addiction transitions occur during the year of the move, and the year of the move is not included in our model moments. In this extension, we assume that addiction transitions do occur during the year of the move, and we assume that these transitions occur according to the place-based addiction parameters of the mover's origin. We continue, however, to not include moments from the

year of the move in our estimation. The result from repeating the exercise detailed in Section 6.2 on this specification is presented in row (8) of Appendix Table A.5.

No Selection

We estimate our specification while assuming there is no selection ($s = 0$). This implies that the distribution of person-specific addiction parameters for each move cohort are identical as for the non-movers in the cohort's origin location, such that

$$\eta_c = \eta_{o(c)}.$$

The result from repeating the exercise detailed in Section 6.2 on this specification is presented in row (9) of Appendix Table A.5.

Alternative Addict Share

As we detail in Appendix E, we normalize the addiction shares across all non-movers observed in our model such that $\bar{a} = 0.10$, and we define this mean addict share formally as

$$\bar{a} = \sum_{t=2006}^{2019} \sum_{j \in \mathscr{J}} \hat{\bar{a}}_{jt} \frac{\hat{n}_{jt}}{N_{nm}},$$

where $\hat{n}_{cr}$ denotes the number of enrollee-years used to construct the moment $\hat{\bar{a}}_{jt}$, $\bar{a}_{jt}$ denotes the average share of addiction for non-movers in a given state-year, and N_{nm} is the total number of non-mover enrollee-years in the sample. In these alternative specifications, we assume that $\bar{a} = 0.15$ and $\bar{a} = 0.20$. The results from repeating the exercise detailed in Section 6.2 on these specifications are presented in rows (10) and (11) of Appendix Table A.5.

Alternative Policy Windows

As we detail in Section 5.2, we allow for policies to be gradually implemented over time through implementation parameters $\zeta^{\pi}_{jt}, \zeta^{\gamma}_{jt} \in [0,1]$, where $\zeta^{\pi}_{jt}, \zeta^{\gamma}_{jt} = 0$ for $t < t^*_j$ and for states that never enact policies, and 1 represents full implementation. In our baseline model, we allow for an implementation window of five years. We assume that policies are fully implemented after five years. In these alternative specifications, however, we vary the length of this window. In particular, we consider a model where policies are fully implemented after three years, and a model where we make no assumption on when policies are fully implemented. The results from repeating the exercise detailed in Section 6.2 on these specifications are presented in rows (12) and (13) of Appendix Table A.5.

Appendix H: Measuring Geographic Characteristics

We construct measures of commuting zone and state-level characteristics—used in Figures 9, Figures A.12, and Appendix Table A.4—from several sources. We briefly describe the measures and their sources here.

Demographics

We use the 2000 Census to construct measures of demography, educational attainment[45], and English proficiency at the commuting zone and the state-level. All of these demographics are population-wide measures.

Health

For the share of the population with obesity or diabetes, we use the 2010 Behavioral Risk Factor Surveillance System (BRFSS), a yearly state-based telephone survey of the non-institutional adult population.[46] The data are compiled and maintained by the CDC Division of Diabetes Translation.

Economic Factors

Share uninsured is the estimated share of the population without insurance, taken from the Census Bureau's 2011 Small Area Health Insurance Estimates.[47] Share with some college is the percent of adults over 25 with some post-secondary education in, and we measure this in the 2000 Census. The unemployment rate is the percent of the population age 16+ unemployed but seeking work in 2012, published by the Bureau of Labor Statistics.[48] China Shock is Autor, Dorn, and Hanson's (2013) measure of exposure to Chinese imports per worker. The variable is defined at the commuting zone (CZ), which are aggregates of counties, and uses variation in commuting zone employment share in industries that are more or less affected by Chinese imports. Specifically, the China Shock is defined by:

$$\Delta IPW_{uit} = \sum_j \frac{L_{ijt}}{L_{ujt}} \frac{\Delta M_{ujt}}{L_{it-1}}$$

where i indexes CZs, j indexes industries, and t indexes time periods. L_{ijt} is CZ i's employment in industry j, L_{ujt} is the total U.S. employment in that industry, and L_{it} is the total employment in CZ i, all in time period $t = 2000$. ΔM_{ucjt} is the realized imports from China to the US in industry j from time periods t (2000) to $t+1$ (2007). Autor, Dorn, and Hanson (2013) explains this variable in more detail. In essence, the variable measures the changes in Chinese imports exposure per worker in each CZ, where imports are apportioned to the region according to its share of national industry employment. We assign each county the China shock associated with that county's CZ.

Finally, employment share in manufacturing is defined in 2000 and derived from the County Business Patterns Data,[49] and median household income is measured using the 2000 Census.[50]

[45] We define the share with some college or higher is measured as the percent of adults over 25 with some post-secondary education.

[46] The share obese is defined as the share of adults in each county that report a BMI greater than or equal to 30; the share diabetic is defined as the share of adults who self-reported being diagnosed with diabetes in 2010.

[47] These data can be downloaded at https://www.census.gov/data/datasets/time-series/demo/sahie/estimates-acs.html.

[48] These data can be downloaded at https://www.bls.gov/lau/#tables.

[49] These data can be downloaded at https://www.census.gov/programs-surveys/cbp/data/datasets.html.

[50] These data can be downloaded at https://www.census.gov/programs-surveys/saipe.html.

Healthcare Characteristics

"MED after surgery" refers to the average morphine equivalent dose (MED) of the prescriptions patients filled in the two weeks after a set of common general surgical procedures. We construct this measure ourselves from the 2006-2019 Medicare claims data for patients (both SSDI and over-65 eligibles) who had not filled an opioid prescription in the 1-30 days prior to the surgery. Our choice of procedures follows Hill et al. (2017), who documented wide variation in opioid prescriptions following several common outpatient surgical procedures at an academic medical center. To identify these surgeries in the claims data we use CPT codes identified by Scully et al. (2018). The surgeries we study are therefore limited to the overlap of procedures studied in the two papers, and consist of laparoscopic cholecystectomy, laparoscopic inguinal hernia repair, and open inguinal hernia repair. Under the (reasonable) assumption that patients do not demand unnecessary surgeries in order to obtain opioids, this measure provides us with a proxy for opioid supply in common surgical settings.

"PCP per capita" measures the number of primary care physicians (PCPs) practicing in the county as a share of the county's total population in 2011, taken from the Area Health Resources Files from information provided by the American Medical Association Masterfile.[51]

"ACSC Rate" is the rate of hospitalizations for ambulatory-care sensitive conditions (ACSC) per Medicare enrollee reported in the 2011 Dartmouth Atlas.[52] ACSCs are preventable conditions (with good outpatient care); high ACSC rates are generally associated with poor healthcare quality (Billings 2003).

"Medicare Spending per Capita" is the total Part A and Part B Medicare spending per Medicare enrollee at the county level, reported by the 2011 Dartmouth Atlas of Health Care.[53]

[51] These data can be downloaded at https://datawarehouse.hrsa.gov/data/datadownload.aspx#MainContent_ctl00_gvDD_lbl_dd_topic_ttl_0.

[52] These data can be downloaded at http://www.countyhealthrankings.org/.

[53] These data can be downloaded at http://www.countyhealthrankings.org/.

Appendix I: Appendix Figures and Tables

Figure A.1: Age-Adjusted Opioid Overdose Death Rates

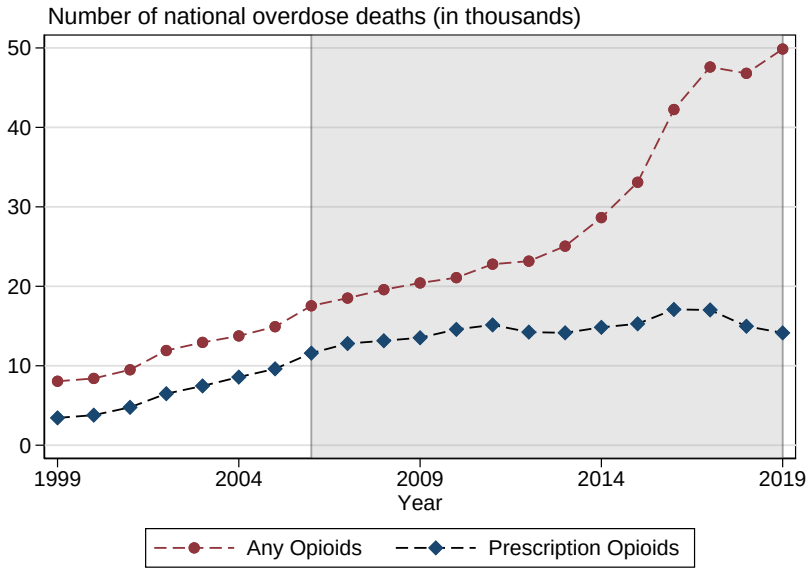

Notes: Figure presents the number of national opioid overdose deaths and the subset of those which involved prescription opioids. The years of our sample – 2006 through 2019 – are shaded. The data was directly obtained from the CDC WONDER Multiple Cause of Death (MCD) Files (National Center for Health Statistics 2021).

Figure A.2: Share of Claims in Destination State by Relative Year

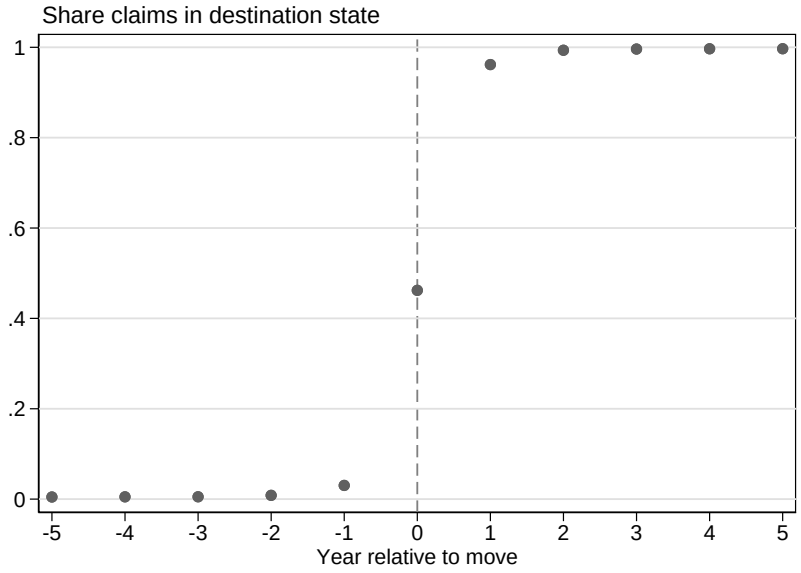

Notes: Figure presents the average fraction of claims in a mover's destination state (as a fraction of claims in either the origin or destination state) by year relative to move. Observations are at the mover-year level. The figure shows a sharp change in the year of the move, with only a small share of claims in the destination pre-move or in the origin post-move. The claim share in the year of the move (relative year 0) is close to 0.5, consistent with moves being roughly uniform throughout the year.

Figure A.3: Opioid-Related Policy Event Studies

(a) Must-Access PDMPs

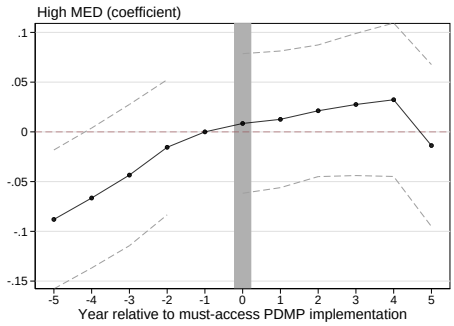

(b) Pill Mill Laws

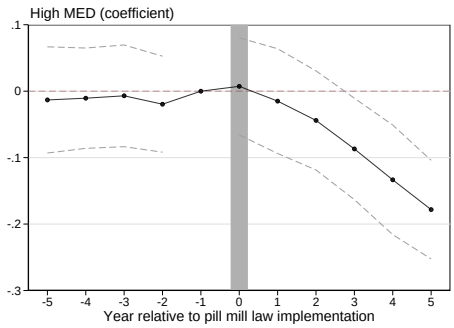

(c) Opioid Prescribing Limits

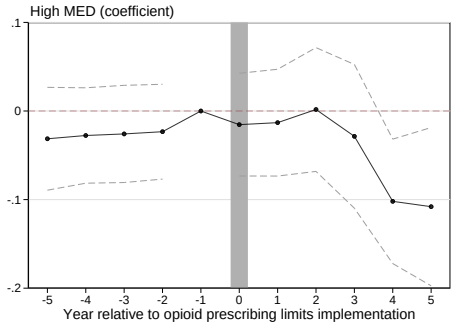

Notes: Figure shows estimates of $\rho_{r(i,t)}$ from equation (6) as estimated using non-movers. Each observation is constructed at the state-year-demographic cell level, and the regression is weighted such that each state-year-demographic cell observation receives weight proportional to the corresponding number of non-mover years observed in the sample. The outcome is risky use rate for each state-year-demographic cell. There are $N = 20,094,127$ enrollee-years in the estimation sample. Dashed lines show 95 percent confidence intervals constructed through 50 iterations of the Bayesian bootstrap procedure (Rubin 1981), which smooths bootstrap samples through reweighting rather than resampling observations. In particular, we draw fifty sets of weights from a gamma distribution for each state-year-demographic cell, and we re-scale the weight given to each observation by these draws. We then re-estimate $\rho_{r(i,t)}$ for each set of drawn weights and compute the standard error across draws. The timing of policies used to calculate relative years is given by Appendix Table A.2, while more details about the specification and the policies are presented in Appendix B.

Figure A.4: Estimating the Effects of Opioid-Related Policies

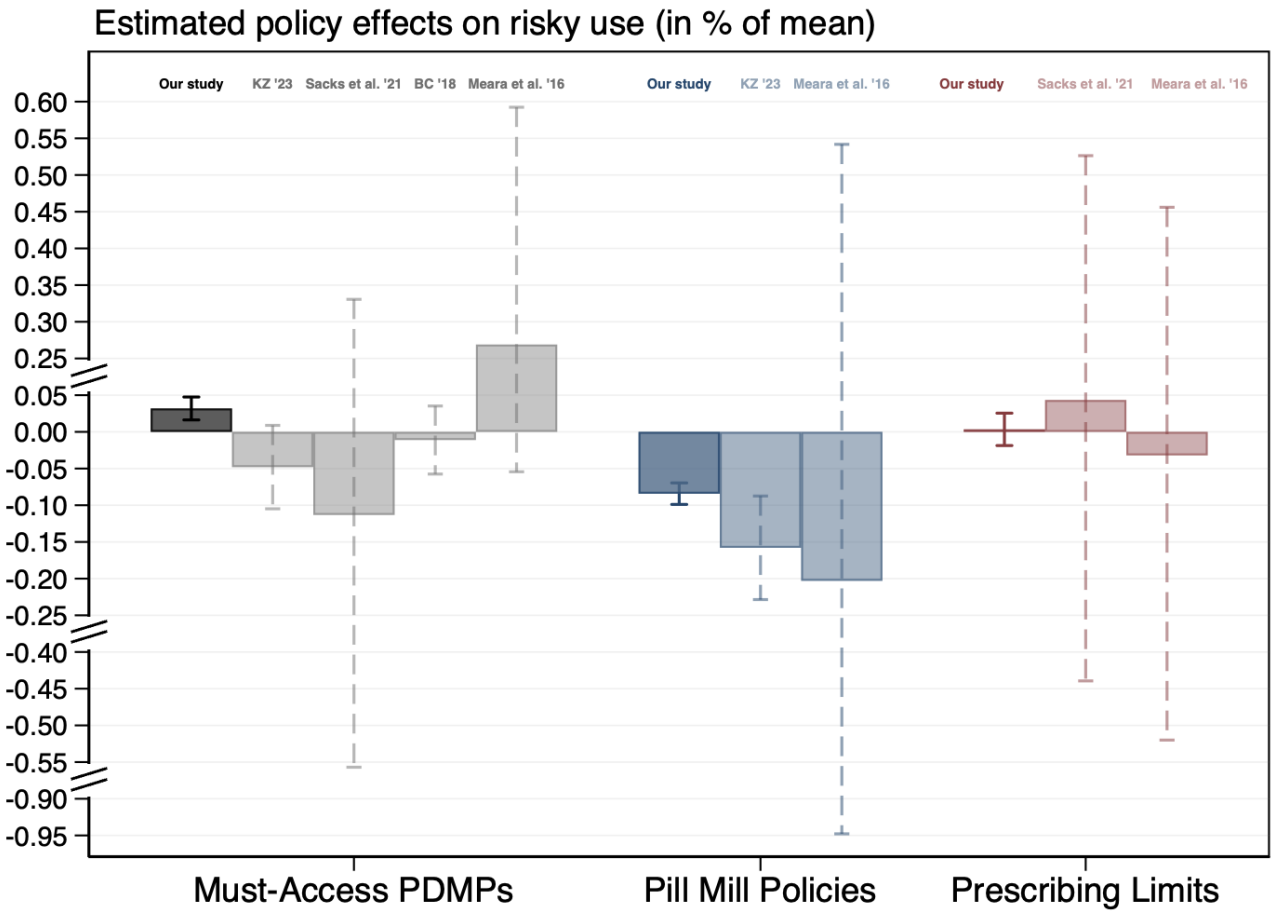

Notes: Figure shows the estimated effects, as a share of the mean, for opioid-related policies from various studies. For four of the studies—Meara et al. (2016),Buchmueller and Carey (2018), Sacks et al. (2021), and our own—we provide their estimates for the policy effects on the probability of opioid prescribing equivalent to over 120 morphine-equivalent doses (MED) in a period. This is simply our baseline measure of risky prescription opioid use, as we discuss in Appendix A. Kaestner and Ziedan (2023) do not study this outcome, so we present their estimated policy effects on changes in opioid sales per capita (as measured in aggregate MED sold per state resident divided by 1,000). We describe each of these studies as well as our own analysis on these policies in Appendix B. Dashed lines show 95 percent confidence intervals taken from each study, and axis breaks on the y-axis are used due to the wide confidence intervals on some of the results. In our own studies, standard errors are constructed through 50 iterations of the Bayesian bootstrap procedure (Rubin 1981), and further detail on the implementation is provided in the exhibit notes of our corresponding opioid-related policy event studies in Appendix Figure A.3.

Figure A.5: Pain Clinic Laws: Legal Passage and Subsequent Amendments

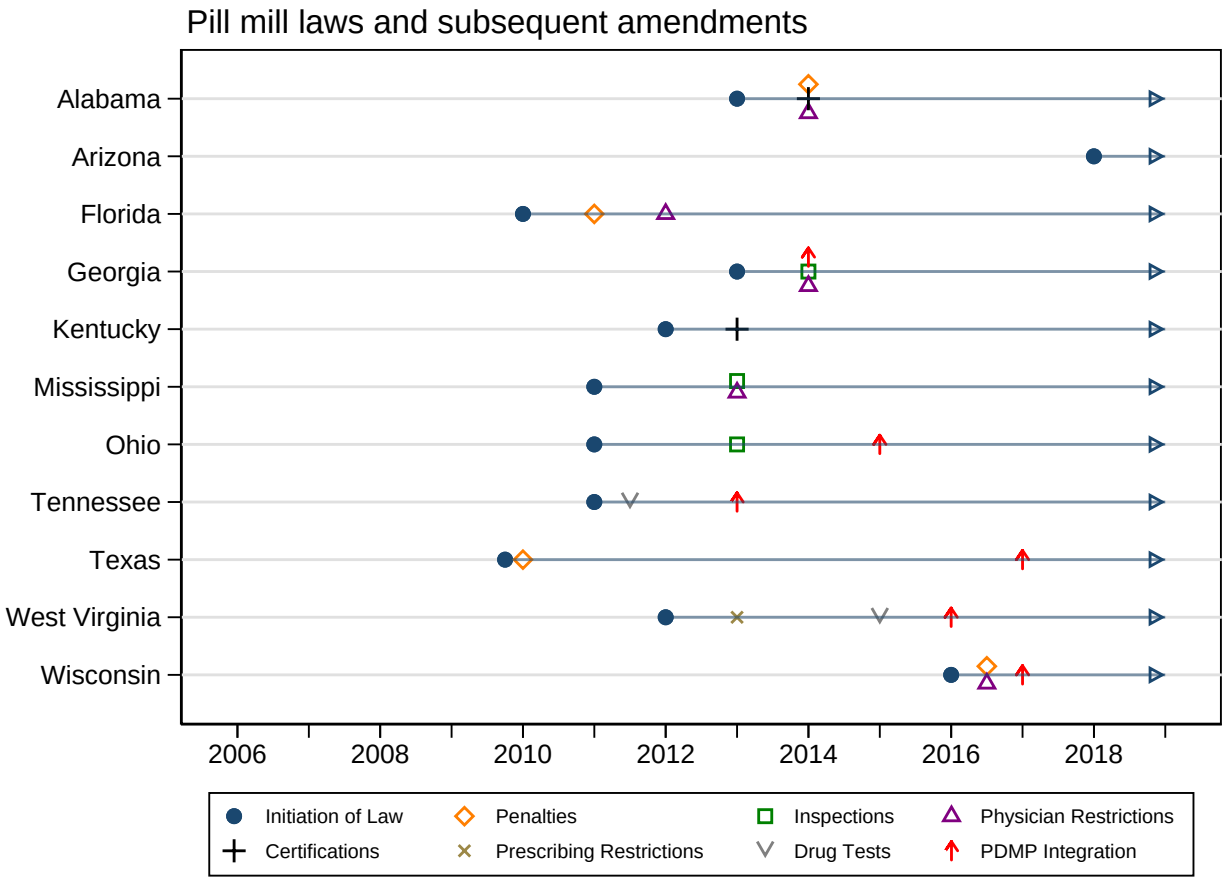

Notes: Figure shows the timing of the passage of pill mill laws for states in our sample as well as the timing of subsequent amendments to these laws in seven categories. Color-coded symbols represent the types of amendments that were made, which include (i) explicit penalties for non-compliance, (ii) explicit inspection of pain clinics to ensure compliance, (iii) restrictions on pain clinic physicians (e.g., no prior felonies, no restrictions on their professional licenses, required training), (iv) certification requirements, (v) restrictions on which types of medical professionals are allowed to prescribe opioids, (vi) requirements for drug testing, and (vii) requirements for prescribers to check PDMPs. We consulted the Prescription Drug Abuse Policy System (PDAPS) and legal databases to measure the timing of pill mill laws and these amendments. More details about pill mill policies are presented in Appendix B.

Figure A.6: Event Studies - Moves Up vs. Moves Down

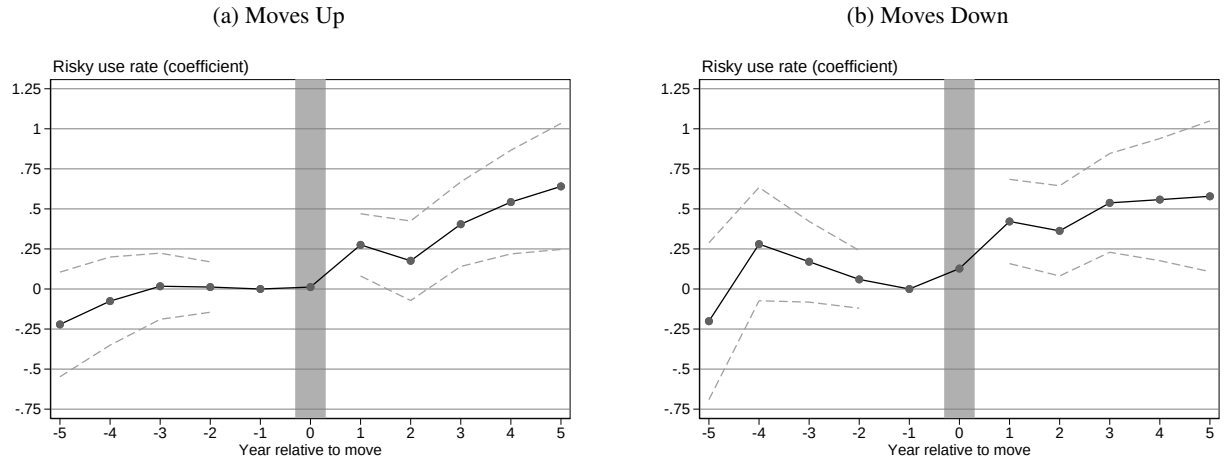

Notes: Figure shows estimates from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort. In Panel A, we restrict our sample to movers who moved to a location with a higher risky use rate (N = 407,211 enrollee-years) while in Panel B, we restrict our sample to movers who moved to a location with a lower risky use rate (N = 318,935 enrollee-years). We define "higher" and "lower" risky use locations by comparing the average rates of risky use among non-movers in each mover's origin and destination during the year of their move. The dashed lines are upper and lower bounds of the 95% confidence interval. Confidence intervals are computed using 50 iterations of the Bayesian bootstrap. Note that our definition of $\hat{T}_{cr}$ in Section 4.1 implies that the coefficient in relative year -1 is zero by construction. The estimates shown in this figure (as well as standard errors) are presented in Appendix Table A.13.

Figure A.7: Event Studies – No Individual Fixed Effects

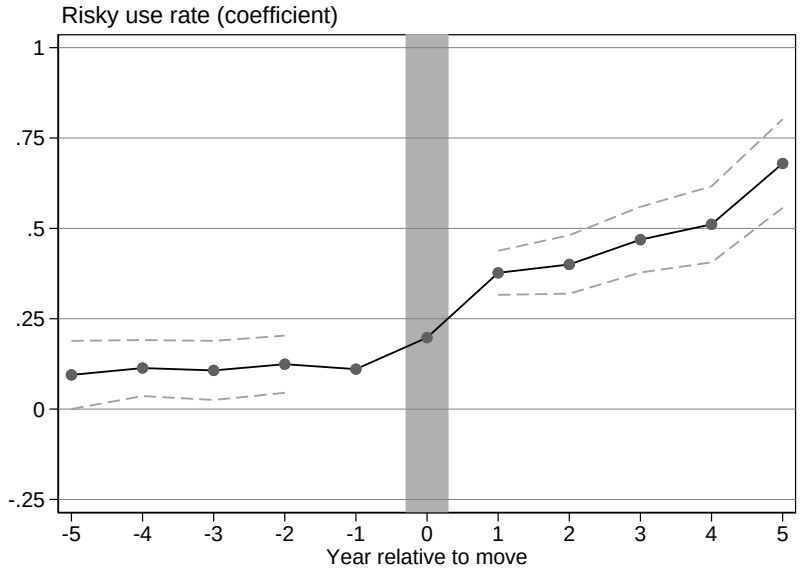

Notes: Figure shows estimates from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort and omitting individual fixed effects (i.e., estimating the model in levels rather than in differences). The dashed lines are upper and lower bounds of the 95% confidence interval. Confidence intervals are computed using 50 iterations of the Bayesian bootstrap. Note that, while our definition of $\hat{T}_{cr}$ in Section 4.1 implies that the coefficient in relative year -1 is zero by construction for event studies when the model is estimated in differences, there is no such mechanical implication for relative year -1 here. This specification is discussed in more detail in Appendix C, and sample sizes are given in Appendix Table A.12.

Figure A.8: Model Moment Fit

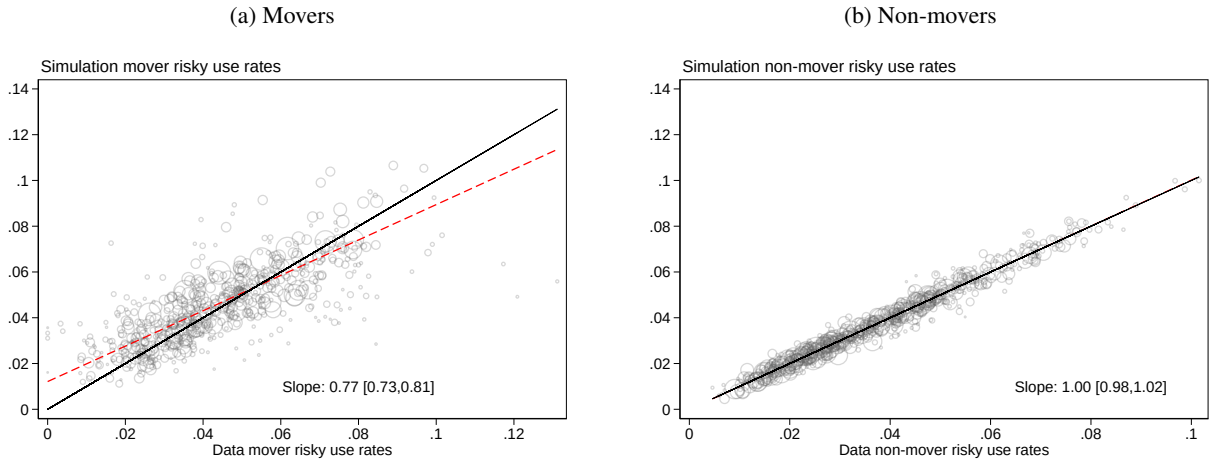

Notes: Figure compares mover (Panel A) and non-mover (Panel B) mean risky opioid use rates predicted from our simulations to those in the Medicare data. Each moment is a state and calendar year combination. In each panel, the y-axis corresponds to the model predicted mean risky use rate, and the solid line in each figure shows the 45 degree line. The size of each observation is proportional to the number of enrollee-years used in constructing the moment. The dashed line shows the line of best fit, using weighted least-squares with weights corresponding to the number of enrollee-years used in constructing the moment. The slope is presented in the bottom left-hand corner, with 95% confidence intervals presented in brackets.

Figure A.9: Simulated vs. Observed State-Level Risky Use Rates

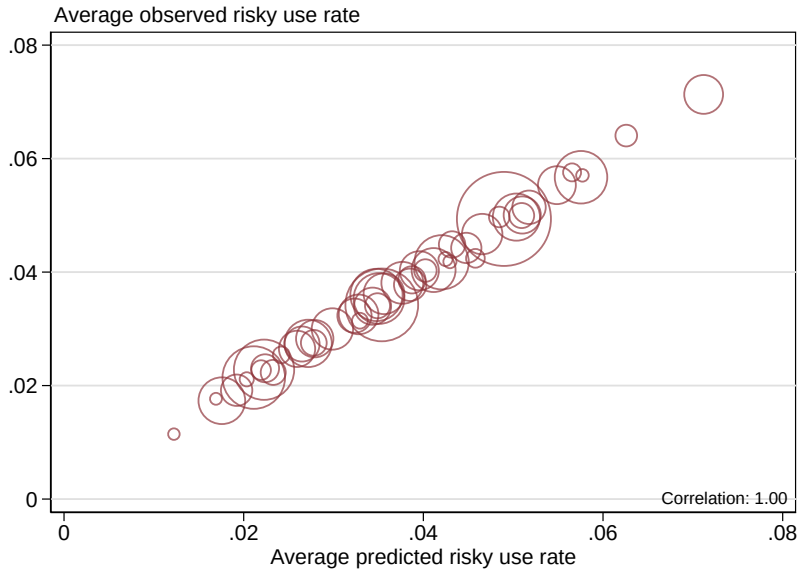

Notes: Figure presents a scatterplot of average state-level observed risky use rates against average state-level predicted risky use rates. Average risky use rates for each state are computed as an average of all non-mover years ($N = 20,094,127$ enrollee-years). The weighted correlation coefficient is presented in the bottom right corner, where weights are given by the number of mover-year observations in our estimation sample where a mover was observed in the state. Markers are scaled accordingly.

Figure A.10: Time Series of Average Risky Opioid Use Rates

(a) Medicare Data

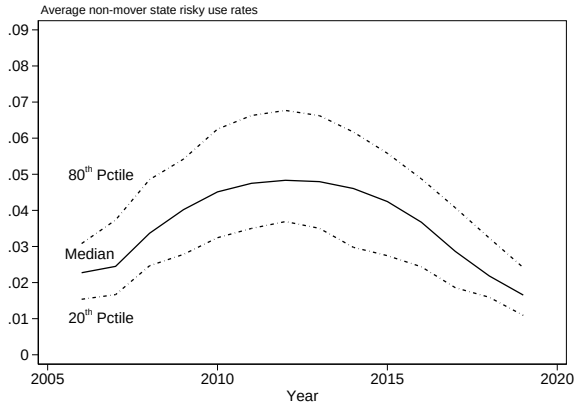

(b) Model Predictions

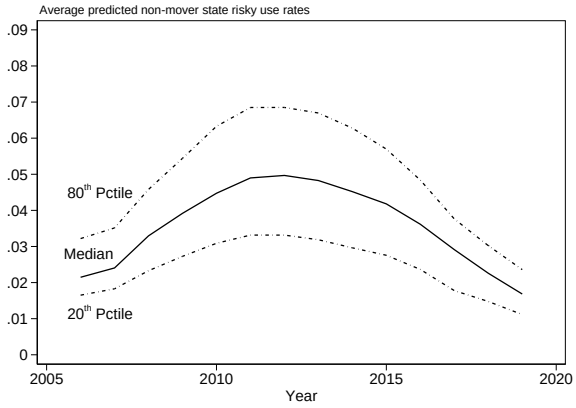

Notes: Figures show time series of average risky use rates among non-movers in the 20th, 50th, and 80th percentile states of the Medicare data (Panel A) and of the model predictions in each year (Panel B). Percentiles are calculated annually, so each line may correspond to different states in each year. Panel A shows results from the Medicare data, while Panel B shows predicted non-mover risky use rates from parameters estimated in our state-level movers only model.

Figure A.11: Pill Mill Policy Event Study in Model and Data

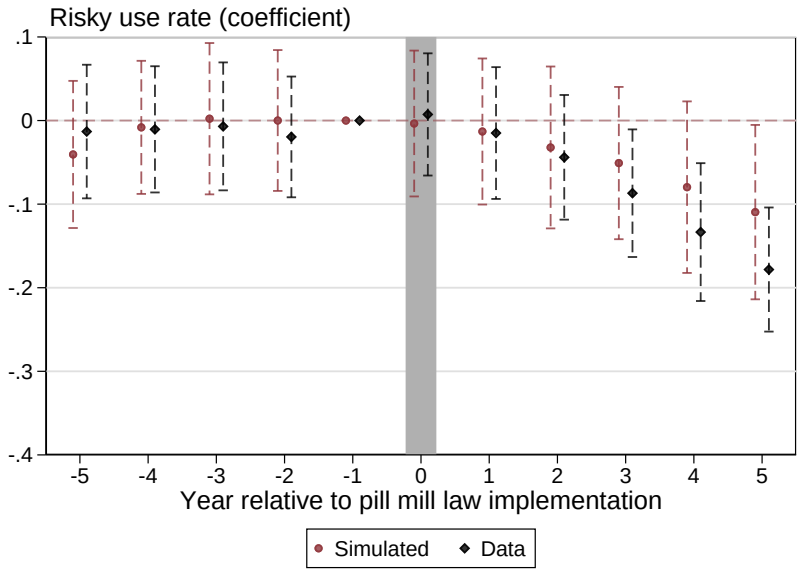

Notes: Figure shows estimates of $\rho_{r(i,t)}$ from equation (8) as estimated on the simulated data and the Medicare data, using non-movers. The simulated data is constructed at the enrollee-year level by using our estimated parameters to simulate individual-level addiction and risky use outcomes for enrollee-years in the baseline sample. The data coefficients are equivalent to Appendix Figure A.3, Panel B where each observation is constructed at the state-year-cell level, and the regression is weighted such that each observation receives weight proportional to the corresponding number of non-mover years observed in the sample. Our model predictions, however, do not vary among demographic cells, so the specification is equivalent to re-writing equation (8) at the state-year level. Dashed lines show 95 percent confidence intervals constructed through 50 iterations of the Bayesian bootstrap. The timing of policies used to calculate relative years is given by Appendix Table A.2, while more details about the specification and the policies are presented in Appendix B.

Figure A.12: State-Level Correlates of Risky Use and Model Parameters

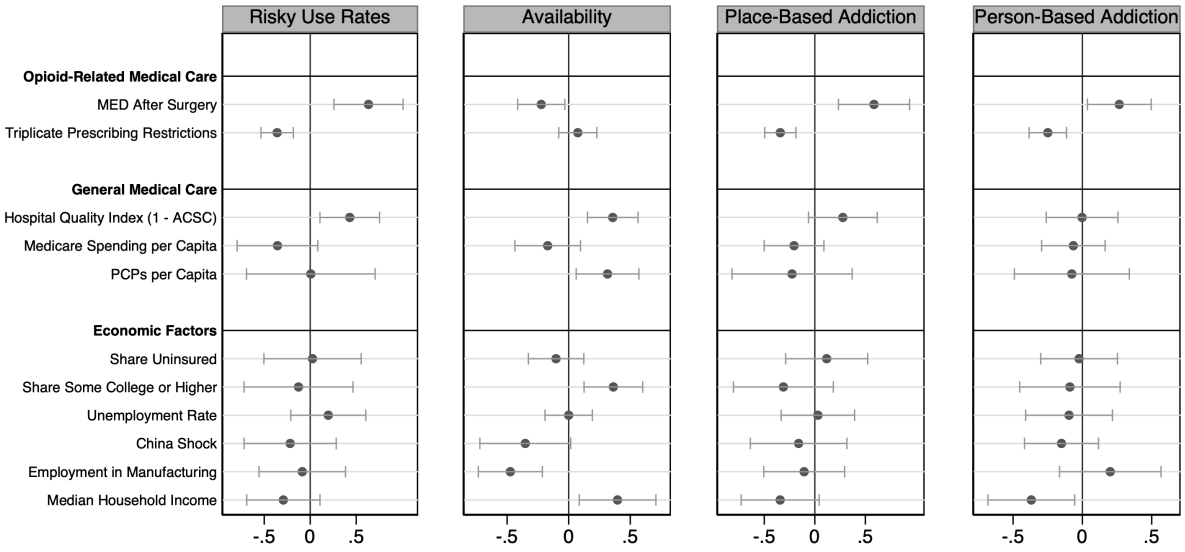

Notes: Figure presents coefficients from bivariate regressions using state-level average characteristics on our baseline model parameters. Horizontal bars show the 95 percent confidence intervals, calculated with heteroskedasticity-robust standard errors. All outcomes and covariates are normalized to be in units of standard deviations. The first panel from the left presents coefficients using the average risky use of the state in our sample as the outcome. The second panel from the left presents coefficients using the estimated availability parameter (γ_j) of that state. The third panel from the left presents the steady state share of addiction in the state implied by using variation in the place-based addiction parameters (π_j) while holding person-based addiction parameters to their median value among all states. Finally, the right-most panel presents the steady state share of addiction in the state implied by using variation in the person-based addiction parameters (η_i) while holding place-based addiction parameters to their median value among all states. All demographic and economic outcomes are measured in the 2000 Census. "MED After Surgery" refers to the average morphine equivalent dose (MED) of the prescriptions patients filled in the two weeks following a set of common surgical procedures among individuals with traditional fee for service coverage in our sample. We measure hospital quality through the "ACSC Rate", or the rate of hospitalizations for ambulatory-care sensitive conditions (ACSC). High ACSC rates are considered a measure of poor healthcare quality, and so we define quality by the rate of avoided hospitalizations for these conditions (1 - ACSC). We present detailed definitions and data sources for these measures in Appendix H.

Figure A.13: Policy Effect Decomposition

(a) No Implementation Delay

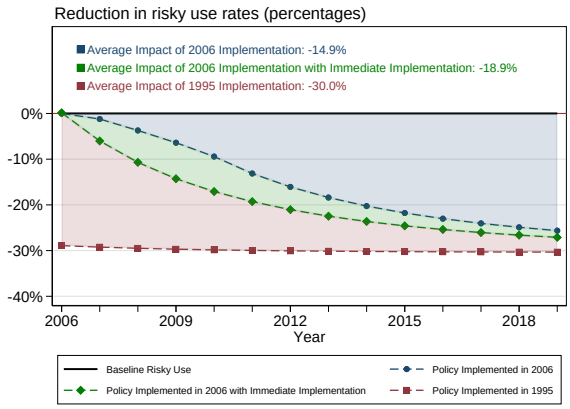

(b) Addiction Stock

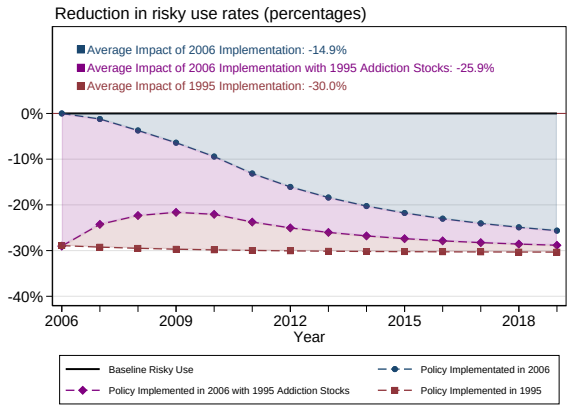

Notes: Figure shows the estimated simulated impact of pill mill laws in the 11 eventually-treated states under various counterfactual implementation timings. These implementation counterfactuals maintain the gradual phase-in of policy effects over five years, as described in Section 5.2, unless noted as "with immediate implementation." Both panels show the reduction in the simulated risky use rates from 2006 to 2019 under a 2006 implementation (blue) and a 1995 implementation (red). Panel A further shows a 2006 implementation in which the policy is fully implemented immediately (green), while Panel B shows an implementation that keeps the gradual policy implementation but where all states start with the (lower) addiction shares in 2006 that they would've had under the earliest 1995 policy implementation. Average impacts are calculated by taking the average reduction across all calendar years between 2006 and 2019.

Figure A.14: Time Trends in Opioid Prescribing Rates

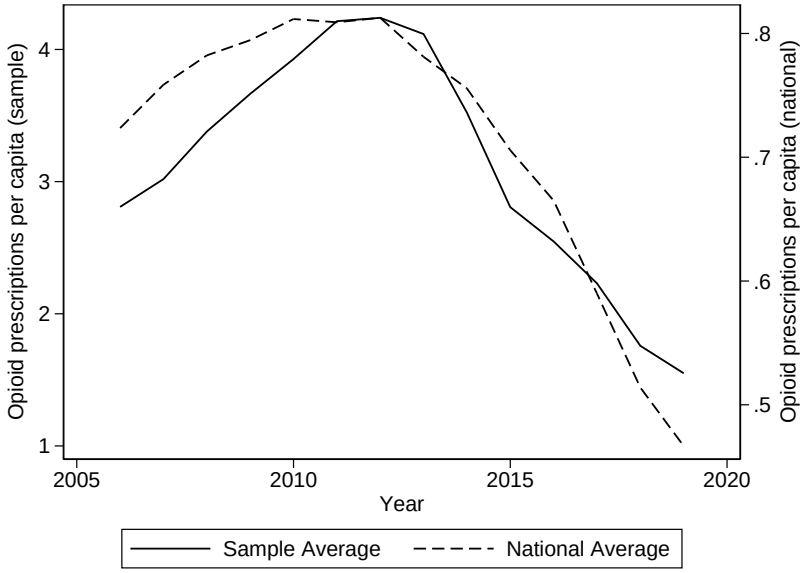

Notes: Figure presents averages over time for 2006 - 2019 of the number of opioid prescriptions per capita among non-movers and from national QuintilesIMS data. Sample prescriptions per capita are plotted on the left-hand vertical axis, and national prescriptions are plotted on the right. Both prescription rates are calculated on a January-December calendar year.

Figure A.15: Event Studies - Various Geographies

(a) Baseline Geography (State)

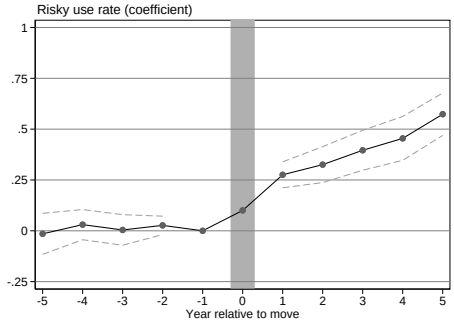

(b) Commuting Zone

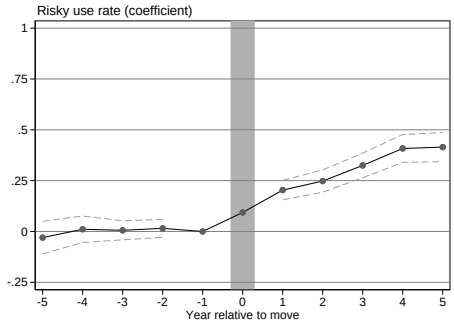

(c) County

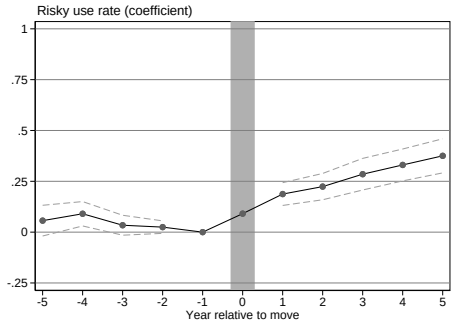

Notes: Figure shows estimates from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort. The dashed lines are upper and lower bounds of the 95% confidence interval. Confidence intervals are computed using 50 iterations of the Bayesian bootstrap. Note that our definition of $\hat{T}_{cr}$ in Section 4.1 implies that the coefficient in relative year -1 is zero by construction. Panel A is our baseline specification, Panel B uses commuting zones as the moving geography, and Panel C uses counties as the moving geography. These specifications are discussed in Appendix C, and sample sizes are given in Appendix Table A.9.

Figure A.16: Event Studies - Other Outcomes

(a) High MED (Baseline)

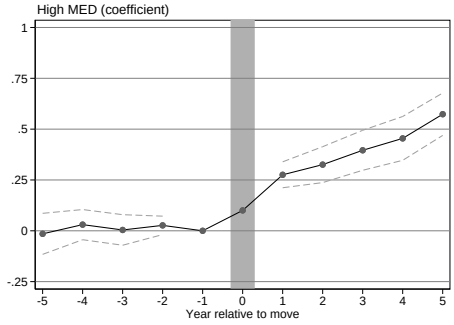

(b) Many Prescribers

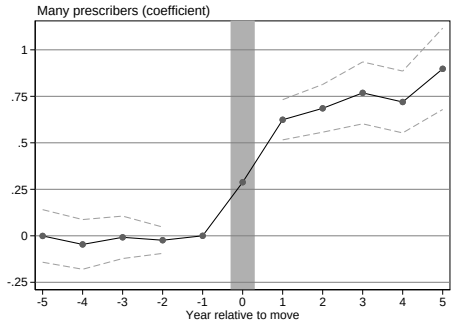

(c) Overlapping Prescriptions

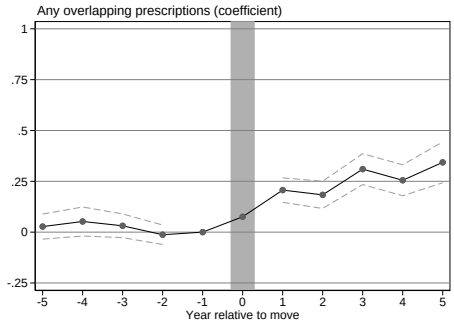

Notes: Figure shows estimates from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort, and y_{ir} (when used to construct $\hat{T}_{cr}$) is defined as the indicated opioid use measure for individual i in relative year r. The dashed lines are upper and lower bounds of the 95% confidence interval. Confidence intervals are computed using 50 iterations of the Bayesian bootstrap. Panel A reports results using the baseline measure (high morphine equivalent doses or "High MED") while Panels B and C report the results with alternate risky use measures of Many Prescribers and Overlapping Prescriptions. These risky use measures are discussed in Appendix A and sample sizes are given in Appendix Table A.11.

Figure A.17: Event Studies: Daily MED and Alternative Thresholds

(a) Daily MED

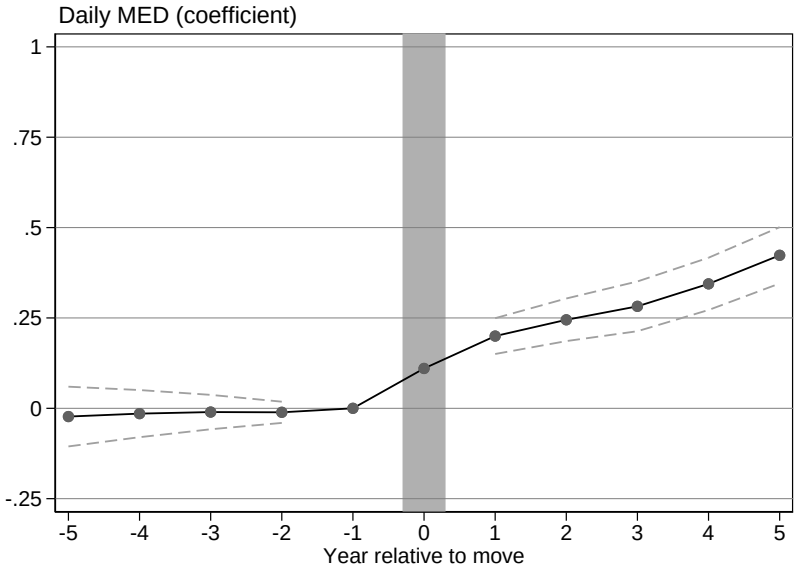

(b) 60 MED Threshold

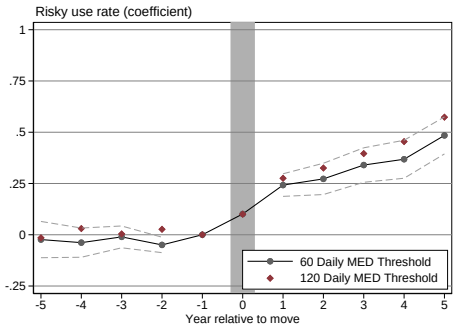

(c) 80 MED Threshold

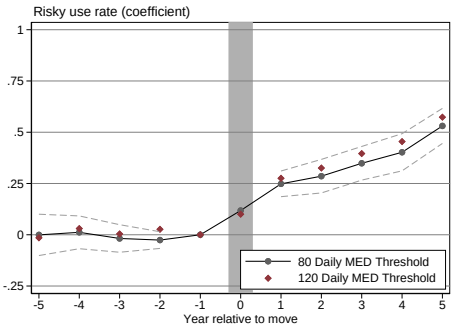

(d) 100 MED Threshold

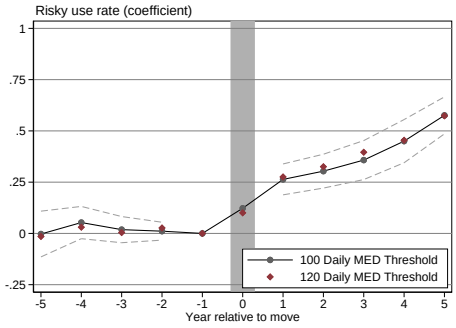

(e) 140 MED Threshold

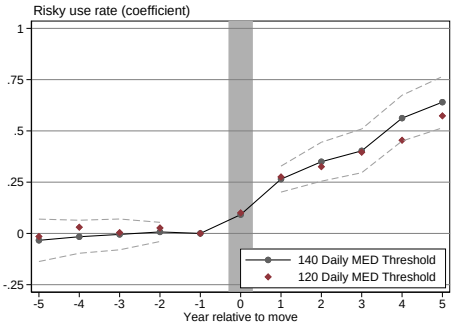

(f) 160 MED Threshold

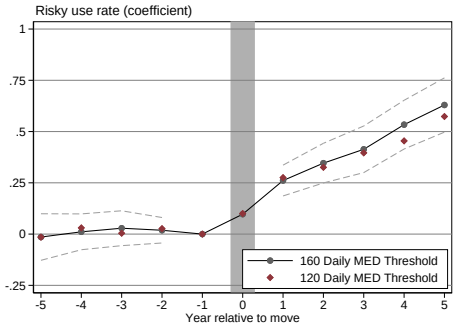

(g) 180 MED Threshold

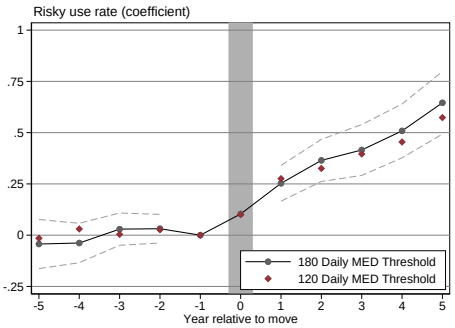

Notes: Figure shows estimates from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort, and the outcomes are defined by alternative functions of the average daily MED prescribed to individuals. Panel A uses the average daily MED over a year, while Panels B - G show alternative thresholds for defining risky use. For Panels B - G, red diamonds are identical to Figure 4 and plot our estimates using our baseline definition of risky use. The dashed lines are upper and lower bounds of the 95% confidence interval. Confidence intervals are computed using 50 iterations of the Bayesian bootstrap.

Figure A.18: Event Studies – Alternative Specifications

(a) No Relative Year Fixed Effects

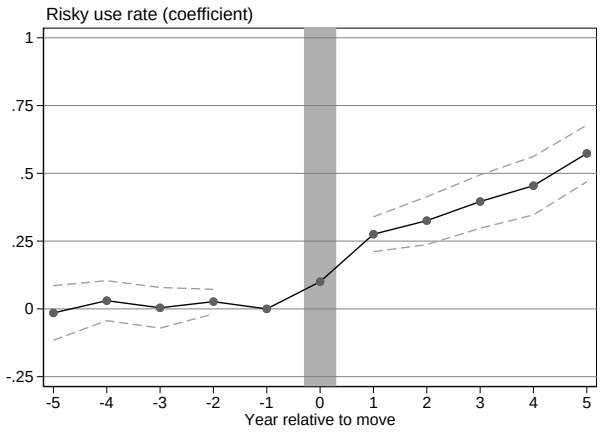

(b) No Covariates (X)

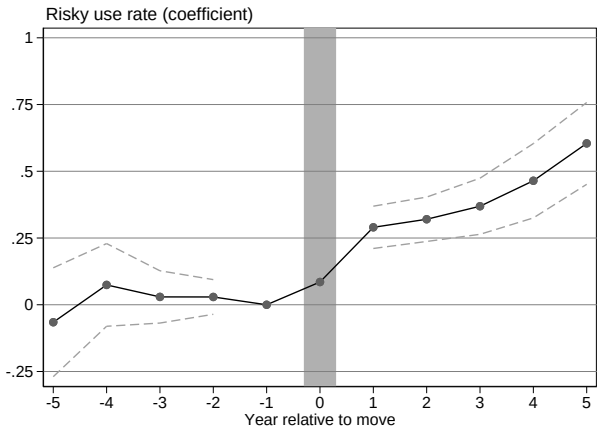

Notes: Figure shows estimates from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort. The dashed lines are upper and lower bounds of the 95% confidence interval. Confidence intervals are computed using 50 iterations of the Bayesian bootstrap. Each panel shows a single deviation from our baseline specification. Panel A omits the relative year fixed effects (ρ_r), and Panel B omits the covariates for five-year age bin, race and gender. These specifications are discussed in more detail in Appendix C, and sample sizes are given in Appendix Table A.12.

Figure A.19: Event Studies – Balanced Panel

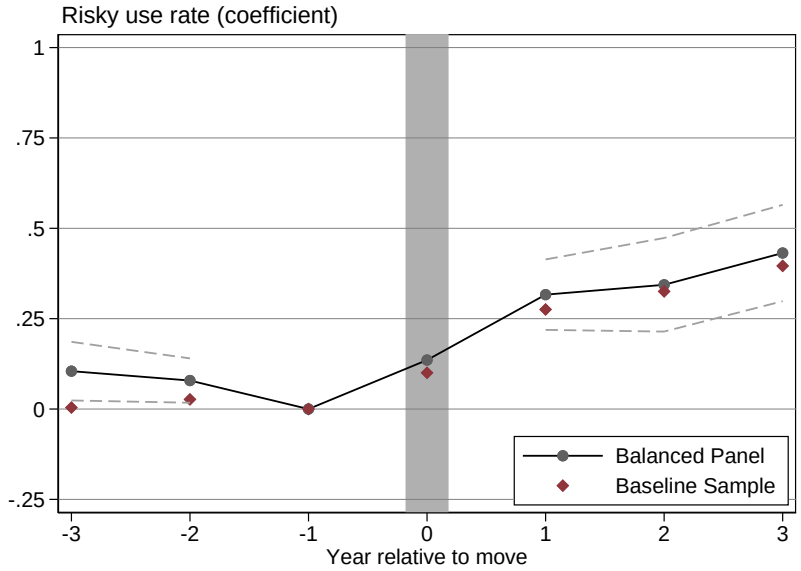

Notes: Gray dots in the figure show estimates from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—for a balanced panel of individuals with observed Part D coverage three years before and three years after the move, where the weights on the regression are given by the number of movers observed in each cohort. The dashed lines are upper and lower bounds of the 95% confidence interval. Red diamonds are identical to Figure 4 and plot our estimates using our baseline sample. Confidence intervals are computed using 50 iterations of the Bayesian bootstrap. Note that our definition of $\hat{T}_{cr}$ in Section 4.1 implies that the coefficient in relative year -1 is zero by construction. The sample consists of movers who are observed for all three years before and after their move (N = 292,087 enrollee-years).

Figure A.20: Event Studies – Sample Restrictions

(a) Always 12 Months of Part D Coverage

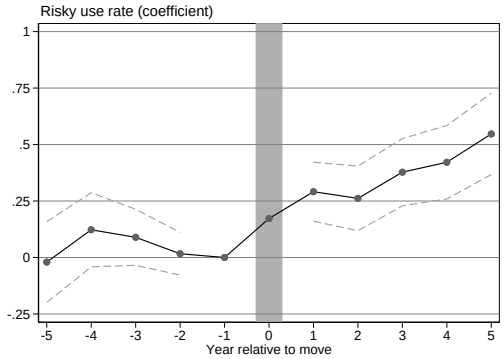

(b) Alive All Years

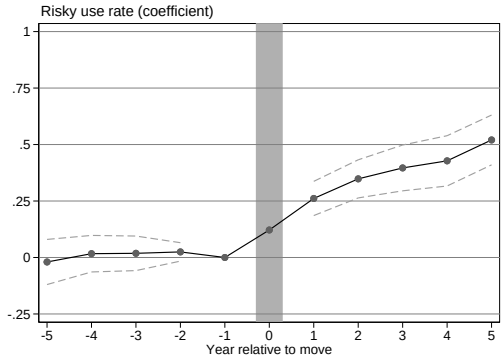

(c) No Moves to/from FL/TX/CA

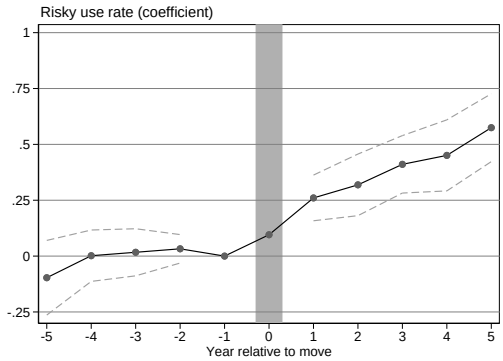

(d) No Cancer/Hospice Patients

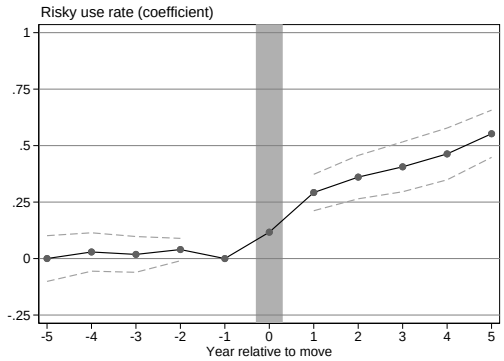

Notes: Figure shows estimates from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort. Panel A restricts to movers with 12 months of Part D coverage for each year they are observed in the sample, and we drop movers who are ever observed with less than a full year of coverage; Panel B restricts to all movers who did not die from 2006 to 2019; Panel C removes all moves to or from Florida, Texas, and California from the sample; and Panel D removes all cancer or hospice patients. These specifications are discussed in Appendix C, and sample sizes are given in Appendix Table A.14. The dashed lines are upper and lower bounds of the 95% confidence interval. Confidence intervals are computed using 50 iterations of the Bayesian bootstrap. Note that our definition of $\hat{T}_{cr}$ in Section 4.1 implies that the coefficient in relative year -1 is zero by construction.

Figure A.21: Event Studies - Selection on Part D Extensive Margin

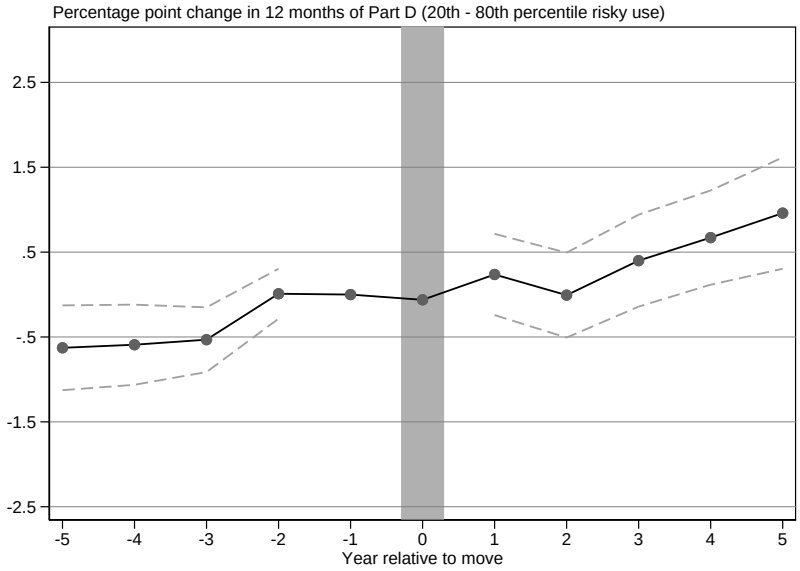

Notes: Figure shows estimates from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort, and y_{ir} (when used to construct $\hat{T}_{cr}$) is defined as having twelve months of Part D coverage during the year. The dashed lines are upper and lower bounds of the 95% confidence interval. Confidence intervals are computed using 50 iterations of the Bayesian bootstrap. Note that our definition of $\hat{T}_{cr}$ in Section 4.1 implies that the coefficient in relative year -1 is zero by construction. The y-axis is scaled to represent a 20th percentile to 80th percentile move in the rate of risky opioid use. The sample is all mover-years, including mover-years without 12 complete months of Part D coverage ($N = 1,096,097$ enrollee-years).

Figure A.22: Commuting Zone Estimation Sample

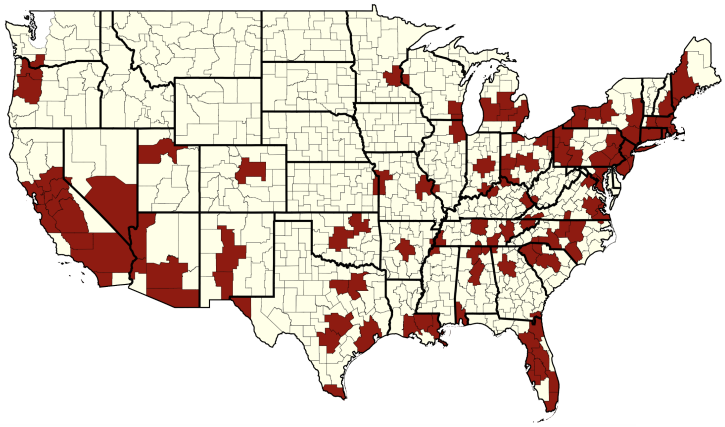

Notes: Figure highlights the 91 commuting zones used for a commuting-zone level model estimation detailed inAppendix C. Thick black lines denote state borders, while thin gray lines denote commuting zone boundaries. Highlighted red regions denote which commuting zones were included.

Figure A.23: Commuting Zone Model Fit

(a) Movers

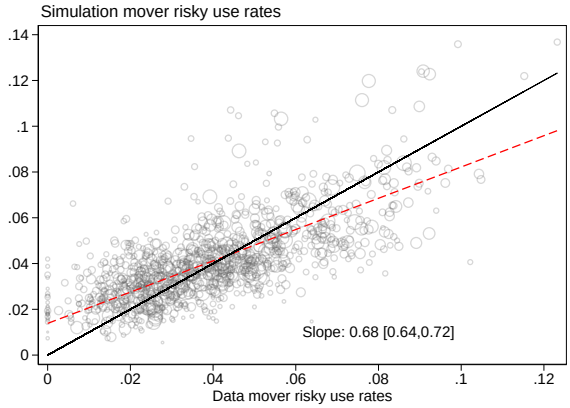

(b) Non-movers

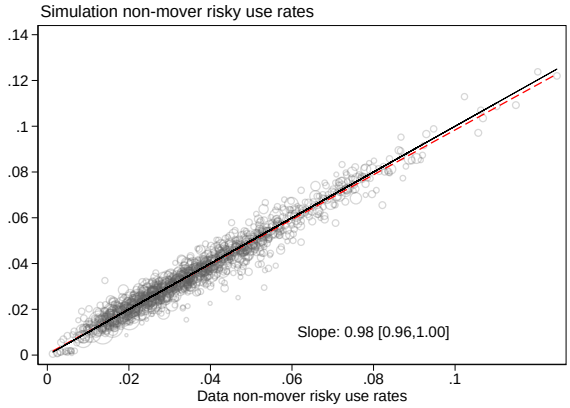

Notes: Figure compares mover (Panel A) and non-mover (Panel B) mean risky opioid use rates predicted from our commuting-zone level model simulations to those in the Medicare data. Each moment is a commuting zone and calendar year combination. In each panel, the y-axis corresponds to the model predicted mean risky use rate, and the solid line in each figure shows the 45 degree line. The size of each observation is proportional to the number of enrollee-years used in constructing the moment. The dashed line shows the line of best fit, using weighted least-squares with weights corresponding to the number of enrollee-years used in constructing the moment. The slope is presented in the bottom left-hand corner, with 95% confidence intervals presented in brackets. More details on the commuting zone sample and estimation are provided in Appendix C.

Figure A.24: Medicare Elderly Sample Event Study

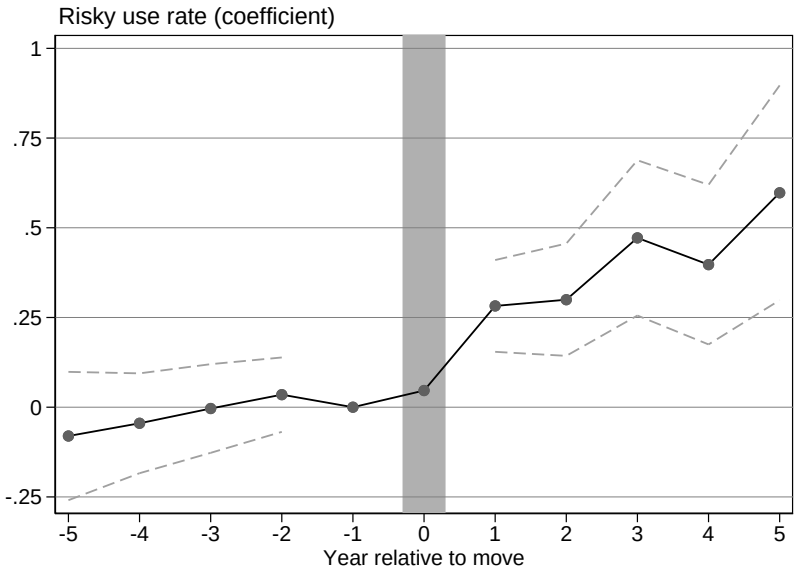

Notes: Figure shows estimates from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort, and the sample used is the Medicare elderly sample described in Appendix C. The dashed lines are upper and lower bounds of the 95% confidence interval. Confidence intervals are computed using 50 iterations of the Bayesian bootstrap. Note that our definition of $\hat{T}_{cr}$ in Section 4.1 implies that the coefficient in relative year -1 is zero by construction. The sample is all mover-years in the Medicare elderly sample ($N = 2,062,303$ mover-years).

Figure A.25: Medicare Elderly Event Studies - Naives and Prior Users

(a) Naives

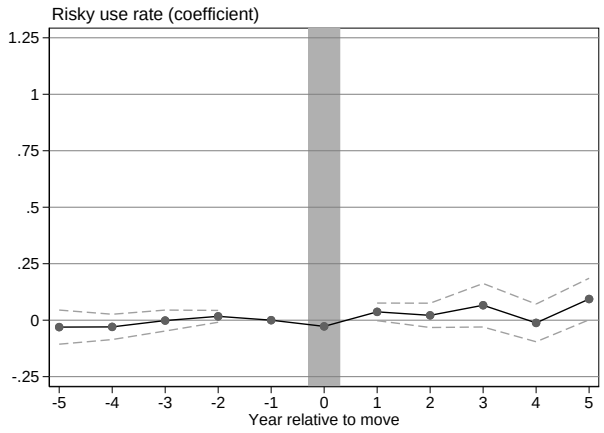

(b) Prior Users

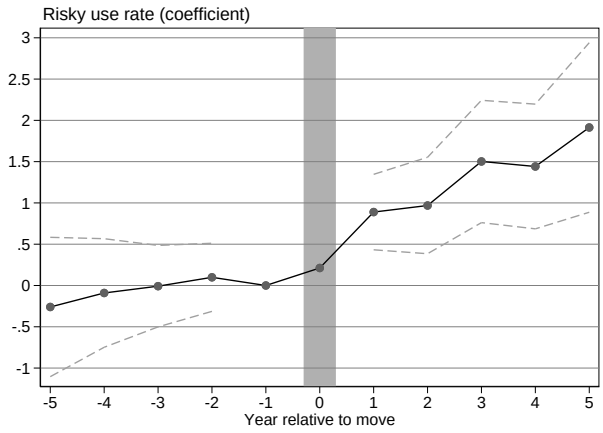

Notes: Figure shows estimates from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort. We show these estimates separately for opioid naive ("naive") and prior users. "Naives" are those with no opioid use in relative year -1, while prior users filled at least one opioid prescription in that year. For each subsample, we extend the vector of observables used for matching to include opioid use in the calendar year corresponding to the year before moving. The dashed lines are upper and lower bounds of the 95% confidence interval. Confidence intervals are computed using 50 iterations of the Bayesian bootstrap. Note that our definition of $\hat{T}_{cr}$ in Section 4.1 implies that the coefficient in relative year -1 is zero by construction. The sample sizes are 1,579,471 mover-years (naives) and 482,674 mover-years (prior users).

Figure A.26: Geographic Variation in Medicare Elderly Model Parameters

(a) Availability (γ_j)

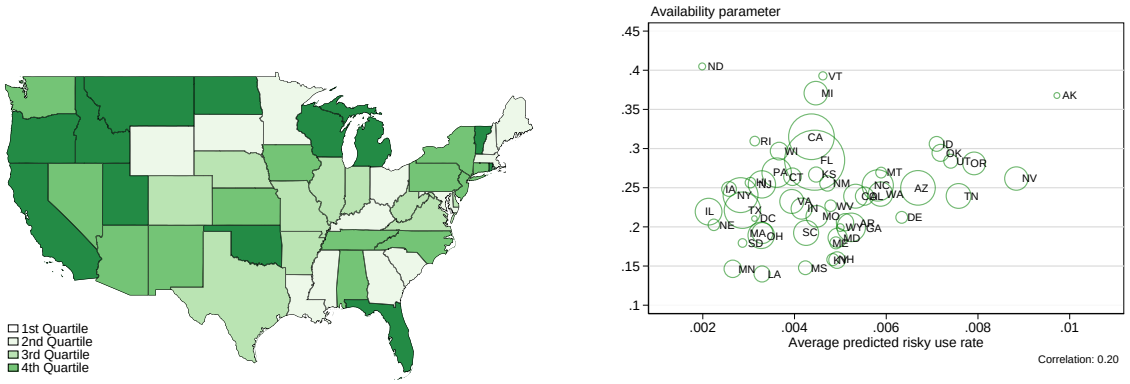

(b) Person-Based Addiction Transitions (η_i)

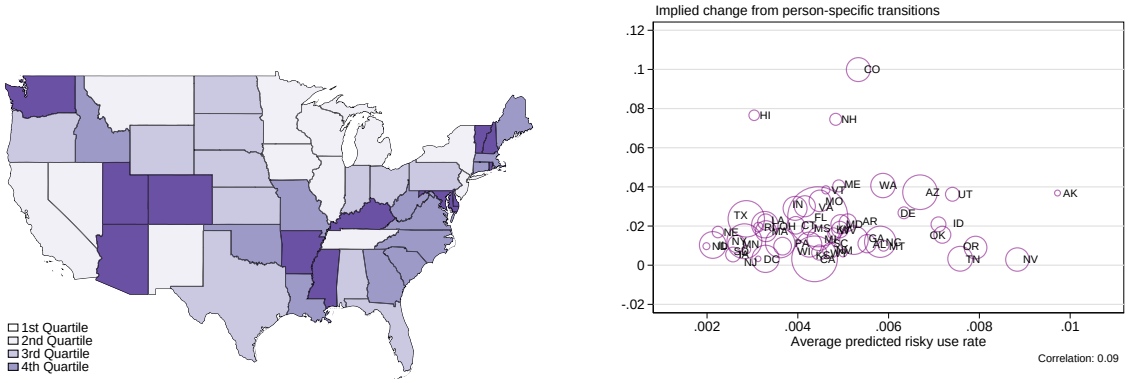

(c) Place-Based Addiction Transitions (π_j)

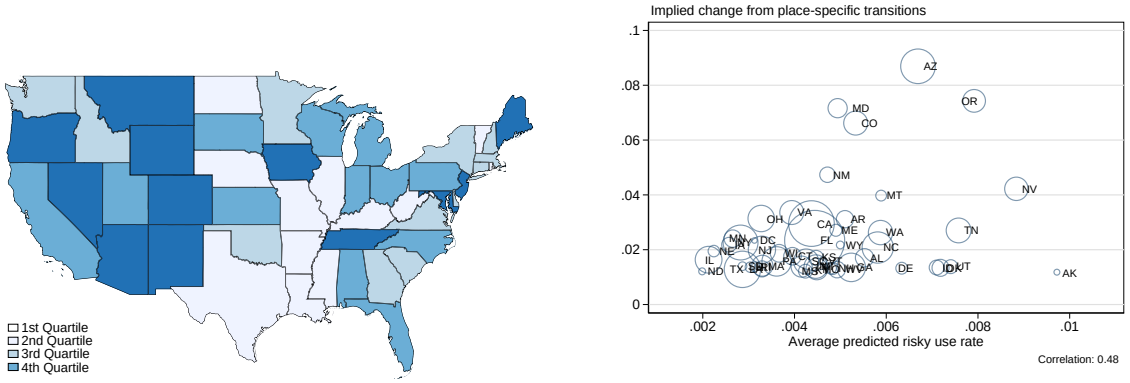

Notes: Figure presents three panels with both maps and scatterplots of parameter estimates for various channels in our model estimates from the Medicare elderly sample described in Appendix C. Panel A presents the place-specific availability parameter, while Panels B and C present person-specific and place-specific factors affecting addiction transitions respectively. The weighted correlation coefficient is presented in the bottom right corner, where weights are given by the number of mover-year observations in the estimation sample where a mover was observed in the state. Markers are scaled accordingly. The magnitude of addiction transition parameters is measured by the change in shares addicted that would result after one year if all states shared an initial share of 0.01 and differential addiction transitions occurred only according to variation in the relevant set of addiction transition parameters. The implied one-year change from place/person specific addiction transitions specifically computes $\Delta\bar{a} = \bar{a}_0(1-\pi^- - \eta^-) + (1-\bar{a}_0)(\pi^+ + \eta^+) - \bar{a}_0$, where $\bar{a}_0$ is set to the 0.01 and the other set of parameters (place vs. person) are held to their median values.

Figure A.27: Moment Fit: Medicare Elderly Sample

(a) Movers

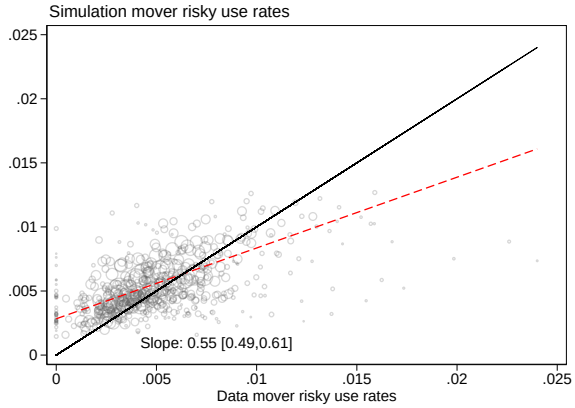

(b) Non-movers

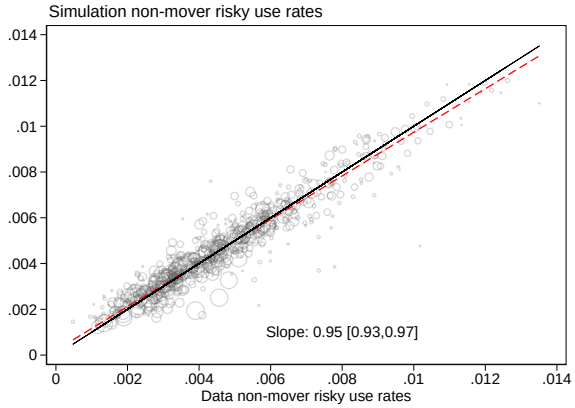

Notes: Figure compares mover (Panel A) and non-mover (Panel B) mean risky opioid use rates predicted from our Medicare elderly model simulations to those in the Medicare data. Each moment is a state and calendar year combination. In each panel, the y-axis corresponds to the predicted mean of the calendar year, and the solid line in each figure shows the 45 degree line. The size of each observation is proportional to the number of enrollee-years used in constructing the moment. The dashed line shows the line of best fit, using weighted least-squares with weights corresponding to the number of enrollee-years used in constructing the moment. The slope is presented in the bottom left-hand corner, with 95% confidence intervals presented in brackets. More details on the Medicare elderly sample are provided in Appendix C.

Figure A.28: Event Study in Simulation vs. Data: Medicare Elderly Sample

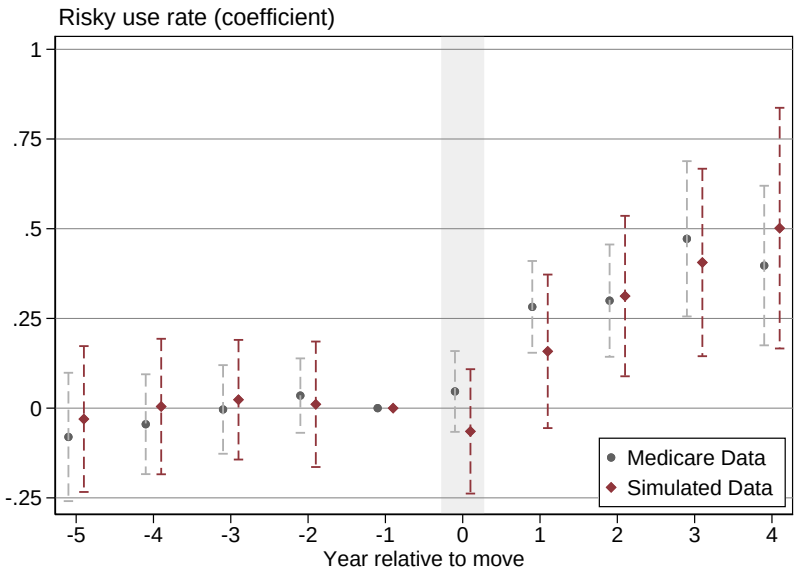

Notes: Figure shows estimates from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—as estimated on simulated data and the Medicare data using our Medicare elderly sample construction. The simulated data is constructed at the enrollee-year level by using our estimated parameters to simulate individual-level addiction and risky use outcomes for enrollee-years in the Medicare elderly sample. More details on the Medicare elderly sample are provided in Appendix C. The weights on the regression are given by the number of movers observed in each cohort. Red diamonds show the aggregate coefficients estimated from our simulated event study, where the yearly rates of risky opioid use among movers and non-movers are simulated according to the model parameterizations described in Section 5 and confidence intervals are calculated across fifty simulations. Gray dots and dashed lines are identical to Appendix Figure A.24 and plot our aggregate event study coefficients and standard errors as estimated on the Medicare elderly sample. Our definition of $\hat{T}_{cr}$ in Section 4.1 implies that the coefficient in relative year -1 is zero by construction.

Figure A.29: Visualizing Hypotheses About Relationships with Illicit Opioid Use

(a) Defining the Modeling Space

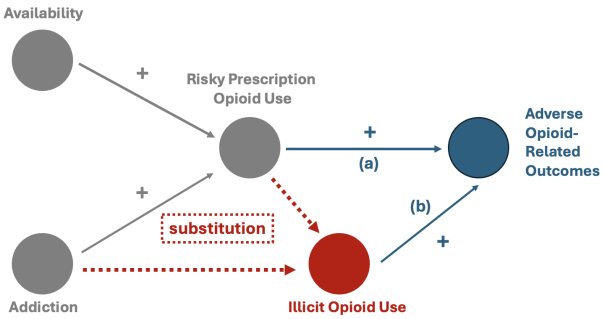

(b) Hypothesis: No Pathway to Illicit Use

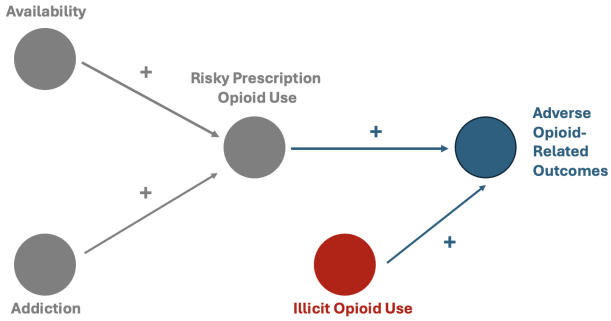

(c) Hypothesis: Substitution Along Addiction Pathway

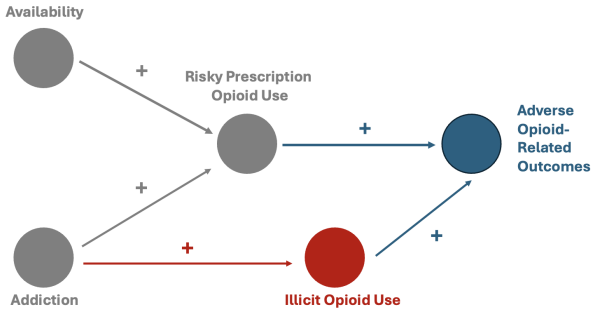

(d) Hypothesis: Substitution Along Both Pathways

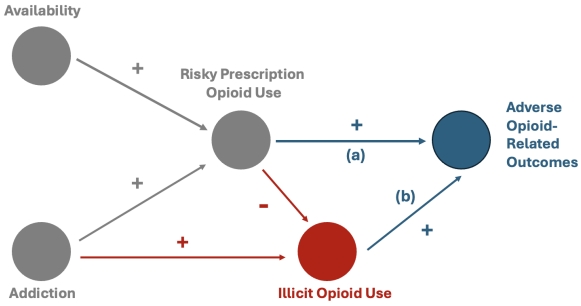

Notes: Figure presents the simple modeling space described in Appendix F where we visualize key hypothesized relationships between availability, addiction, risky prescription opioid use, illicit opioid use, and adverse opioid-related outcomes. Arrows represent directed causal relationships, while positives and negatives indicate the sign of this relationship. We assume that all relationships can be strictly signed. Panel A presents an overview of the limited modeling space that we consider, where availability and addiction positively drive risky use (denoted in gray), while risky prescription opioid use and illicit use drive adverse opioid-related outcomes (denoted in blue). The two dashed red lines denote the subject of the hypotheses in Appendix F. Panel B considers a model in which neither addiction nor risky use directly affect illicit use. Panel C considers a model in which addiction drives both risky prescription opioid use and illicit opioid use, but changes in risky prescription opioid usage do not directly affect illicit opioid usage. Finally, Panel D considers a model in which there is a direct positive relationship between addiction and illicit opioid use and a direct negative relationship between risky prescription opioid use and illicit opioid use. Further assumptions are detailed in Appendix F.

Table A.1: Opioid Prescribing Measures and Adverse Opioid Outcomes

	Measure of Risky Use	Share of All Enrollee-Years	Share Opioid Poisoning Event in Year $t+1$	Enrollee-Years with Opioid Use Disorder in Year $t+1$
(1)	High MED (Baseline)	2.98%	0.87%	3.93%
(2)	Many Prescribers	4.44%	0.85%	3.76%
(3)	Overlapping Prescribers	13.09%	0.57%	2.34%
(4)	Chronic Opioid Use	17.77%	0.46%	1.89%
(5)	Any Opioids	35.94%	0.28%	1.19%
(6)	All Years	100.00%	0.11%	0.53%

Notes: Table shows the incidence of adverse opioid outcomes (opioid poisonings and diagnoses of opioid use disorders) in subsequent enrollee-years when enrollee-years exhibit various measures of risky opioid use and opioid prescribing. We limit the sample to non-mover enrollee years fully covered by traditional Medicare, whom we also observe as fully covered by traditional Medicare in the following year. This allows us to observe their adverse outcomes in year $t+1$. The first column defines the measure of opioid use, where "all years" is the full sample unconditional on any opioid use ($N = 5,930,877$ enrollee-years). The second column presents the share of all enrollee-years in our sample that exhibit the given measure of opioid use; the third and fourth columns present the share of these observations with adverse opioid outcomes in the year following. These various risky opioid use measures are discussed in Appendix A.

Table A.2: Opioid-Related Policy Implementation Timing

State	Must-Access PDMP	Pill Mill Policies	Prescribing Limits
Alabama	–	2013	–
Alaska	2017	–	2017
Arizona	2017	2018	–
Arkansas	2017	–	–
California	2018	–	–
Colorado	–	–	–
Connecticut	2015	–	2016
Delaware	–	–	2017
Florida	–	2010	–
Georgia	2014	2013	–
Hawaii	–	–	2016
Idaho	–	–	–
Illinois	2018	–	2012
Indiana	2014	–	2017
Iowa	–	–	–
Kansas	–	–	–
Kentucky	2012	2012	2017
Louisiana	–	–	2017
Maine	–	–	–
Maryland	–	–	2017
Massachusetts	2014	–	–
Michigan	–	–	2016
Minnesota	2017	–	2017
Mississippi	–	2011	–
Missouri	–	–	–
Montana	–	–	–
Nebraska	–	–	–
Nevada	2015	–	2017
New Hampshire	2016	–	2017
New Jersey	2015	–	2017
New Mexico	2012	–	–
New York	2013	–	2016
North Carolina	–	–	2018
North Dakota	–	–	–
Ohio	–	2011	2017
Oklahoma	–	–	–
Oregon	–	–	–
Pennsylvania	2017	–	2017
Rhode Island	2016	–	2017
South Carolina	2017	–	2007
South Dakota	–	–	–
Tennessee	2013	2011	2013
Texas	2019	2010	–
Utah	2017	–	2017
Vermont	2015	–	2017
Virginia	2015	–	2017
Washington	–	–	–
West Virginia	–	2012	–
Wisconsin	2017	2016	–
Wyoming	–	–	–

Notes: Table presents the year of "implementation" that we use for our policy analysis throughout the paper (2006 - 2019). This excludes implementation dates before or after this window. We attempt to follow the legal databases, the Prescription Drug Abuse Policy System (PDAPS), and literature analyzing the effects of opioid-related policies that were enacted during our study period (Buchmueller and Carey 2018; Sacks et al. 2021; Kaestner and Ziedan 2023). We discuss each of these policies—and uncertainty regarding the timing across these sources—in Appendix B.

Table A.3: Estimation Sample Summary Statistics

	Movers	Non-Movers
Sample Used to Construct Moments		
Number of Enrollees	99,729	2,755,447
Number of Enrollee-Years	637,647	20,094,127
Number of Cohorts	16,612	–
Average Risky Use Rate	4.1%	3.5%
Moments Targeted in Model Estimation		
Number of Moments	125,921	714
Median Number of Enrollee-Years per Moment	2	19,624.5
Mean Number of Enrollee-Years per Moment	5.1	28,143.0
Minimum Number of Enrollee-Years per Moment	1	728
Maximum Number of Enrollee-Years per Moment	313	167,054

Notes: Table presents summary statistics on the empirical moments that are targeted in the estimation of the model as well as the sample used to construct those moments. The sample used to construct the moments is simply the sample of all movers presented in Appendix Table 1 column 1, excluding enrollee-years observed during the year of move. The empirical moments are average risky use rates constructed at the cohort – relative-year level, where cohorts are defined by a mover's origin, destination, and move year.

Table A.4: Full Set of CZ-Level Correlates of Risky Use and Model Parameters

(a) Correlates from Main Text

	(1)	(2)	(4)	(5)
Covariate	Risky Use	Availability	Place-Based Addiction	Person-Based Addiction
MED after Surgery	0.403	-0.024	0.358	0.068
	(0.108)	(0.131)	(0.094	(0.105
Triplicate Prescribing Restrictions	-0.418	-0.030	-0.259	-0.212
	(0.080)	(0.087)	(0.072)	(0.087)
Hospital Quality Index (ACSC Rate)	-0.225	-0.287	-0.143	0.065
	(0.103	(0.114)	(0.098)	(0.136)
Medicare Spending per Capita	-0.148	-0.041	-0.172	0.142
	(0.098)	(0.140)	(0.067)	(0.126)
PCPs per Capita	0.044	0.153	-0.111	0.140
	(0.110)	(0.086)	(0.094)	(0.094)
Share Uninsured	-0.011	0.002	0.064	-0.041
	(0.107)	(0.112)	(0.084)	(0.101)
Share with Some College or Higher	-0.100	0.402	-0.230	-0.061
	(0.102)	(0.091)	(0.087)	(0.097)
Unemployment Rate	-0.018	-0.101	-0.008	-0.057
	(0.113)	(0.103)	(0.106)	(0.120)
Exposure to China Shock	-0.073	-0.018	-0.072	-0.017
	(0.099)	(0.109)	(0.079)	(0.099)
Employment in Manufacturing	-0.014	-0.344	0.073	0.015
	(0.123)	(0.107)	(0.103)	(0.112)
Median Household Income	-0.047	0.314	-0.207	-0.034
	(0.106)	(0.112)	(0.100)	(0.112)

(b) Additional Correlates

	(1)	(2)	(4)	(5)
Covariate	Risky Use	Availability	Place-Based Addiction	Person-Based Addiction
Share White	0.345	-0.072	0.207	0.217
	(0.085)	(0.092)	(0.085)	(0.090)
Share Black	-0.211	-0.056	-0.175	0.035
	(0.108	(0.111)	(0.099)	(0.107)
Share Hispanic	-0.286	0.089	-0.142	-0.258
	(0.077)	(0.102)	(0.075)	(0.083)
Share Asian	-0.141	0.115	-0.137	-0.085
	(0.088)	(0.076)	(0.079)	(0.065)
Share Female	-0.243	-0.025	-0.249	0.080
	(0.119)	(0.098)	(0.095)	(0.083)
Share Under 18	-0.276	-0.123	-0.107	-0.291
	(0.112)	(0.090)	(0.105)	(0.104)
Share Over 65	0.254	0.041	0.096	0.232
	(0.101)	(0.100)	(0.082)	(0.122)
Share SSDI in 1990	0.09	-0.412	0.251	0.078
	(0.141)	(0.098)	(0.150)	(0.147)
Share with Poor English	-0.286	0.161	-0.197	-0.193
	(0.078)	(0.100)	(0.067)	(0.082)
Share Obese	0.048	-0.346	0.138	0.084
	(0.105)	(0.109)	(0.092)	(0.100)
Share Diabetic	0.098	-0.338	0.158	0.198
	(0.116)	(0.107)	(0.108)	(0.126)
Share Uninsured	-0.011	0.002	0.064	-0.041
	(0.107)	(0.112)	(0.084)	(0.101)

Notes: Table presents coefficients from bivariate regressions using commuting-zone level average characteristics on risky use and our model parameters and their associated heteroskedasticity-robust standard errors. All outcomes and covariates are normalized to be in units of standard deviations. Column (1) presents coefficients using the average risky use of the commuting zone in our sample as the outcome. Column (2) presents coefficients using the estimated availability parameter (γ_j) of that commuting zone. Column (3) presents the steady state share of addiction in the commuting zone implied by using variation in the place-based parameters (π_j) while holding the person-based addiction parameters to their median value among all CZs. Column (4) presents the steady state share of addiction in the commuting zone implied by using variation in the person-based addiction parameters (η_i) while holding place-based addiction parameters to their median value among all CZs. Panel A shows the correlates presented in Table A.8 in the main text, while Panel B shows additional correlates. All demographic and economic outcomes besides "Share SSDI in 1990" are measured in the 2000 Census. "Share SSDI in 1990" is measured following Cutler and Glaeser (2021). "MED after surgery" refers to the average morphine equivalent dose (MED) of the prescriptions patients filled in the two weeks following a set of common surgical procedures in our sample. "ACSC rate" is the rate of hospitalizations for ambulatory-care sensitive conditions (ACSC); high ACSC rates are considered a measure of poor healthcare quality. We present detailed definitions and data sources for these measures in Appendix H.

Table A.5: Robustness of Counterfactuals to Alternative Models

		Predicted reduction in rate of risky use from a one standard deviation reduction in...				
	Model Specification	Person-based Parameters (η_i)	Place-based Parameters (γ_j, π_j)	Availability (γ_j)	Place-based Addiction Transitions (π_j)	All Addiction Transitions (π_j, η_i)
(1)	Baseline Model	32.6%	41.3%	13.9%	31.9%	51.0%
	Alternative Samples					
(2)	Commuting Zone Sample	36.4%	58.5%	13.8%	51.4%	72.2%
(3)	Medicare Elderly (65+) Sample	36.8%	53.5%	18.0%	43.0%	60.8%
	Alternative Specifications					
(4)	Multiplicative Addiction Interactions($\alpha \to 0$)	58.0%	38.1%	11.2%	30.4%	69.1%
(5)	Minimum Addiction Interactions($\alpha \to -\infty$)	30.8%	45.5%	27.3%	25.1%	40.6%
(6)	Maximum Addiction Interactions($\alpha \to \infty$)	22.5%	31.7%	10.9%	23.2%	41.9%
(7)	Allowing for Risky Use by Non-Addicted Individuals	16.0%	23.0%	9.3%	16.6%	24.9%
(8)	Addiction Transitions During Year of Move	42.4%	47.6%	13.4%	39.3%	67.8%
(9)	No Selection ($s = 0$)	61.2%	47.5%	11.7%	40.5%	74.9%
(10)	Alternative Global Addiction Share ($\bar{a} = 0.15$)	29.3%	38.3%	12.8%	29.3%	46.7%
(11)	Alternative Global Addiction Share ($\bar{a} = 0.20$)	29.5%	36.8%	11.0%	29.2%	49.3%
(12)	Three Years Policy Adoption Window	33.1%	42.9%	14.7%	33.2%	52.5%
(13)	All Years Policy Adoption Window	25.6%	39.8%	14.5%	29.8%	44.9%

Notes: Table presents the effect of counterfactuals for our baseline specification as well as a set of alternative samples and models. Estimates are presented in percent reduction in simulated overall risky use rates from 2006 to 2019 from a one standard deviation reduction in parameters. Each counterfactual policy is to reduce parameters in the relevant channel(s) by a standard deviation, except for the probabilities of transitioning out of addiction, which we increase by a standard deviation. The overall risky use rate is computed by taking a weighted average of simulated risky use over all state-years in our non-mover sample, weighting by the number of non-movers observed in each state-year. Throughout, we maintain our standard parameter bounds of [0,1] by taking the minimum (maximum) of the reduced (increased) parameter and the bound as necessary. Row (1) presents the results from our baseline specification. Rows (2) and (3) present the results from a CZ-level subsample and a Medicare elderly sample described in Appendix C. Rows (4) - (6) examine models where person-based and place-based addiction transition factors interact non-additively. Row (7) examines an alternative model where we allow for risky use by non-addicted individuals at a constant rate, which is an additional parameter that we estimate. Row (8) examines a model where we allow for addiction transitions to occur according to the origin destination parameters in the year of the move. Row (9) estimates our model when we assume that there is no selection ($s = 0$). Rows (10) and (11) examine models which specify different global shares of addiction ($\bar{a} = 0.15$ and $\bar{a} = 0.20$). Finally, rows (12) and (13) allow for different assumptions about when pill mill laws are fully implemented. The three-year policy window refers to a model in which after three years, pill mill policies are fully implemented. The all-years policy window refers to a model in which we never assume pill mill policies are fully implemented (subject to a normalization). We discuss these specifications further in Appendix G.

Table A.6: Yearly Effects of Counterfactual Policies Targeting Risky Use

	Targeting Place-Based Addiction Channel	Targeting Place-Based Availability Channel	Targeting Both Place-Based Channels
2006	0.0%	19.0%	19.0%
	[0.0%, 0.0%]	[10.3%, 27.7%]	[10.3%, 27.7%]
2007	10.5%	17.6%	26.2%
	[3.4%, 17.5%]	[10.8%, 24.4%]	[19.7%, 32.7%]
2008	18.3%	13.7%	29.4%
	[9.8%, 26.7%]	[8.3%, 19.0%]	[22.2%, 36.7%]
2009	24.1%	11.9%	33.2%
	[14.9%, 33.3%]	[7.2%, 16.7%]	[25.0%, 41.3%]
2010	28.6%	10.6%	36.2%
	[18.8%, 38.5%]	[6.4%, 14.9%]	[27.2%, 45.3%]
2011	32.2%	10.0%	39.0%
	[21.5%, 42.9%]	[6.0%, 14.0%]	[29.0%, 48.9%]
2012	35.1%	10.1%	41.6%
	[23.6%, 46.5%]	[6.0%, 14.2%]	[30.9%, 52.3%]
2013	37.4%	10.6%	44.0%
	[25.3%, 49.5%]	[6.3%, 15.0%]	[32.7%, 55.4%]
2014	39.3%	11.6%	46.3%
	[26.7%, 52.0%]	[6.9%, 16.3%]	[34.5%, 58.2%]
2015	40.9%	12.7%	48.4%
	[27.8%, 54.1%]	[7.5%, 17.9%]	[36.2%, 60.7%]
2016	42.3%	14.9%	50.8%
	[28.7%, 55.8%]	[8.8%, 21%]	[38.3%, 63.4%]
2017	43.3%	19.0%	54.1%
	[29.4%, 57.3%]	[11.2%, 26.8%]	[41.3%, 66.8%]
2018	44.2%	24.5%	57.8%
	[30.0%, 58.4%]	[14.4%, 34.6%]	[45.0%, 70.5%]
2019	44.9%	32.8%	62.7%
	[30.4%, 59.3%]	[19.3%, 46.2%]	[50.0%, 75.5%]

Notes: Table reports the yearly percentage reduction in risky use rates driven by the effect of counterfactual policies that lower place-based addiction and place-based availability parameters as described in Section 6.3 and presented in Figure 10. Each counterfactual policy is to reduce parameters in the relevant channel(s) by a standard deviation except for the probabilities of transitioning out of addiction, which we increase by a standard deviation. The yearly risky use rate is computed by taking a weighted average of predicted risky use over all states in a given year in our non-mover sample, weighting by the number of non-movers observed in each state. 95% confidence intervals are computed using 50 iterations of the Bayesian bootstrap. Throughout, we maintain our standard parameter bounds of [0,1] by taking the minimum (maximum) of the reduced (increased) parameter and the bound as necessary.

Table A.7: Morphine Equivalents Conversion Table

Opioid Active Ingredient	Morphine Equivalents per Milligram
Codeine	0.15
Dihydrocodeine	0.25
Fentanyl (transdermal)	2.4
Fentanyl (oral)	0.1
Hydrocodone	1
Hydromorphone	4
Levorphanol	12
Meperidine	0.1
Methadone	4
Morphine	1
Oxycodone	1.5
Oxymorphone	3
Pentazocine	0.3
Propoxyphene	0.6
Tapentadol	0.367

Notes: Table identifies opioid conversion factors used in the construction of our baseline measure of risky use. Adapted from Meara et al. (2016) supplementary material as well as an opioid dose calculator developed by Washington State Agency Medical Directors' Group available at http://www.agencymeddirectors.wa.gov/calculator/dosecalculator.html.

Table A.8: Measures of Hazardous Opioid Prescriptions

	High MED (Baseline)	Many Prescribers	Overlapping Prescriptions	Any Opioid	Chronic Opioid Use
High MED (Baseline)	1.000				
Many Prescribers	0.685	1.000			
Overlapping Prescriptions	0.756	0.765	1.000		
Any Opioids	0.469	0.630	0.907	1.000	
Chronic Opioid Use	0.529	0.595	0.929	0.976	1.000

Notes: Table presents pairwise rank correlations between the baseline measure (high morphine equivalent doses or "High MED"), Many Prescribers, Overlapping Prescribers, Any Opioids, and Chronic Opioid Use variables for 2006 - 2019, averaged across years at the state level and assigning equal weight to each state. These various risky opioid use measures are discussed in Appendix A.

Table A.9: Event Study Coefficients – Various Geographies

	(1)	(2)	(3)
	All	Naive	Prior User
Panel A: Baseline Geography (State)			
Average of 1-5 years pre-move	0.009	-0.008	0.054
	(0.030)	(0.018)	(0.072)
Average of 1-5 years post-move	0.405	0.127	0.710
	(0.047)	(0.025)	(0.096)
Enrollees	99,729	57,334	42,341
Enrollee-years	726,146	412,821	312,866
Panel B: Commuting Zone			
Average of 1-5 years pre-move	0.000	0.001	0.013
	(0.024)	(0.014)	(0.055)
Average of 1-5 years post-move	0.320	0.082	0.586
	(0.031)	(0.019)	(0.077)
Enrollees	126,847	73,877	51,762
Enrollee-years	916,727	529,951	376,868
Panel C: County			
Average of 1-5 years pre-move	0.0411	0.014	0.0842
	(0.022)	(0.013)	(0.053)
Average of 1-5 years post-move	0.280	0.073	0.554
	(0.037)	(0.019)	(0.092)
Enrollees	130,373	75,465	50,461
Enrollee-years	927,813	533,402	358,356

Notes: Table reports the average of coefficients and bootstrapped standard errors (in parentheses) one to five years pre-move and one to five years post-move, as estimated from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort. Estimates are shown for the baseline sample, naive enrollees, and prior users observed in the year before move for whom we are able to observe matched non-movers. Panel A reports estimates for our baseline result, Panel B reports estimates for moves by commuting zone, and Panel C reports estimates for moves by county. Opioid naive enrollees are those with no opioid use in relative year -1, while prior users filled at least one opioid prescription in that year. Standard errors are computed using 50 iterations of the Bayesian bootstrap.

Table A.10: State-Year Level Risky Opioid Use Measure Summary Statistics

	25th Percentile	Median	75th Percentile	Interquartile Range
High MED (Baseline)	0.023	0.035	0.049	0.026
Many Prescribers	0.029	0.048	0.069	0.040
Overlapping Prescriptions	0.096	0.130	0.170	0.074

Notes: Table presents the 25th percentile, median, and 75th percentile of the state-year averages of the indicated opioid prescription measure: our baseline measure (high morphine equivalent doses or "High MED"), Many Prescribers, and Overlapping Prescriptions. These various risky opioid use measures are discussed in Appendix A. Each state-year average is determined by averaging the opioid prescription measure outcome within a year across non-movers in the sample. The final column presents the interquartile range of each opioid prescription measure.

Table A.11: Event Study Coefficients – Other Outcomes

	(1)	(2)	(3)
	All	Naive	Prior User
Panel A: High MED (Baseline)			
Average of 1-5 years pre-move	0.009	-0.008	0.054
	(0.030)	(0.018)	(0.072)
Average of 1-5 years post-move	0.405	0.127	0.710
	(0.047)	(0.025)	(0.096)
Enrollees	99,729	57,334	42,341
Enrollee-years	726,146	412,821	312,866
Panel B: Many Prescribers			
Average of 1-5 years pre-move	-0.016	-0.002	-0.040
	(0.047)	(0.029)	(0.105)
Average of 1-5 years post-move	0.739	0.261	1.257
	(0.080)	(0.049)	(0.157)
Enrollees	99,729	57,334	42,341
Enrollee-years	726,146	412,821	312,866
Panel C: Overlapping Prescriptions			
Average of 1-5 years pre-move	0.020	0.013	0.053
	(0.024)	(0.020)	(0.053)
Average of 1-5 years post-move	0.260	0.145	0.416
	(0.039)	(0.029)	(0.079)
Enrollees	99,729	57,334	42,341
Enrollee-years	726,146	412,821	312,866

Notes: Table reports the average of coefficients and their bootstrapped standard errors (in parentheses) between one to five years pre-move and one to five years post-move as estimated from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort and y_{ir} (when used to construct $\hat{T}_{cr}$) is defined as the indicated opioid use measure for individual i in relative year r. Estimates are shown for the sample of all movers, naive enrollees, and prior users observed in the year before move for whom we are able to observe matched non-movers. Each panel reports results for a different outcome: our baseline measure (high morphine equivalent doses or "High MED"), Many Prescribers, and Overlapping Prescriptions. These risky opioid use measures are discussed in Appendix A. Opioid naive enrollees are those with no opioid use in relative year -1, while prior users filled at least one opioid prescription in that year. Standard errors are computed using 50 iterations of the Bayesian bootstrap.

Table A.12: Event Study Coefficients – Alternative Specifications

	(1)	(2)	(3)
	All	Naive	Prior User
Panel A: Baseline Specification			
Average of 1-5 years pre-move	0.009	-0.008	0.054
	(0.030)	(0.018)	(0.072)
Average of 1-5 years post-move	0.405	0.127	0.710
	(0.047)	(0.025)	(0.096)
Enrollees	99,729	57,334	42,341
Enrollee-years	726,146	412,821	312,866
Panel B: No Individual Fixed Effects			
Average of 1-5 years pre-move	0.105	-0.008	0.216
	(0.041)	(0.018)	(0.093)
Average of 1-5 years post-move	0.4755	0.1271	0.906
	(0.040)	(0.025)	(0.096)
Enrollees	138,115	57,334	42,341
Enrollee-years	878,724	412,821	312,866
Panel C: No Relative Year Fixed Effects			
Average of 1-5 years pre-move	0.009	-0.008	0.054
	(0.030)	(0.018)	(0.072)
Average of 1-5 years post-move	0.405	0.127	0.710
	(0.047)	(0.025)	(0.096)
Enrollees	99,729	57,334	42,341
Enrollee-years	726,146	412,821	312,866
Panel D: No Covariates (X)			
Average of 1-5 years pre-move	0.009	-0.005	0.040
	(0.030)	(0.018)	(0.071)
Average of 1-5 years post-move	0.401	0.124	0.698
	(0.048)	(0.024)	(0.098)
Enrollees	99,782	57,387	42,395
Enrollee-years	726,594	413,270	313,324

Notes: Table reports the average of coefficients and their bootstrapped standard errors (in parentheses) between one to five years pre-move and one to five years post-move as estimated from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort. Estimates are shown for the baseline sample, naive enrollees, and prior users observed in the year before move for whom we are able to observe matched non-movers. Each panel shows a single deviation from our baseline specification (baseline shown in Panel A). Panel B omits individual fixed effects (estimating the model in levels rather than in differences); Panel C omits the relative year fixed effects (ρ_r), and Panel D omits the covariates for five-year age bin, race and gender. Standard errors are computed using 50 iterations of the Bayesian bootstrap.

Table A.13: Event Study Coefficients – Moves Up and Moves Down

	(1)	(2)	(3)
	All	Naive	Prior User
Panel A: Baseline (All Moves)			
Average of 1-5 years pre-move	0.009	-0.008	0.054
	0.030	0.018	0.072
Average of 1-5 years post-move	0.405	0.127	0.710
	0.047	0.025	0.096
Enrollees	99,729	57,334	42,341
Enrollee-years	726,146	412,821	312,866
Panel B: Moves Up			
Average of 1-5 years pre-move	-0.005	0.016	-0.145
	0.053	0.047	0.186
Average of 1-5 years post-move	0.494	0.128	0.735
	0.131	0.071	0.292
Enrollees	55,722	18,150	16,254
Enrollee-years	407,211	121,832	107,909
Panel C: Moves Down			
Average of 1-5 years pre-move	0.010	-0.014	0.145
	0.067	0.054	0.253
Average of 1-5 years post-move	0.397	0.098	0.899
	0.112	0.056	0.305
Enrollees	44,007	14,332	14,290
Enrollee-years	318,935	96,090	94,898

Notes: Table reports the average of coefficients and their bootstrapped standard errors (in parentheses) between one to five years pre-move and one to five years post-move as estimated from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort. Each panel shows a single deviation from our baseline specification (baseline shown in Panel A). The table reports estimates for moves up (Panel B) and moves down (Panel C) based on the difference in the rates of risky opioid use between the origin and destination state in the year of the move. Opioid naive enrollees are those with no opioid use in relative year -1, while prior users filled at least one opioid prescription in that year. For each subsample, we extend the vector of observables used for matching to include opioid use in the calendar year corresponding to the year before moving. Estimates are shown for the baseline sample, naive enrollees, and prior users observed in the year before move for whom we are able to observe matched non-movers. Standard errors are computed using 50 iterations of the Bayesian bootstrap.

Table A.14: Event Study Coefficients – Sample Restrictions

	(1)	(2)	(3)
	All	Naive	Prior User
Panel A: Baseline Sample			
Average of 1-5 years pre-move	0.009	-0.008	0.054
	(0.030)	(0.018)	(0.072)
Average of 1-5 years post-move	0.405	0.127	0.710
	(0.047)	(0.025)	(0.096)
Enrollees	99,729	57,334	42,341
Enrollee-years	726,146	412,821	312,866
Panel B: Part D All Years			
Average of 1-5 years pre-move	0.042	0.007	0.107
	(0.057)	(0.029)	(0.144)
Average of 1-5 years post-move	0.380	0.177	0.623
	(0.078)	(0.042)	(0.164)
Enrollees	22,556	13,277	9,224
Enrollee-years	195,617	113,730	81,460
Panel C: Alive All Years			
Average of 1-5 years pre-move	0.008	-0.016	0.067
	(0.030)	(0.020)	(0.075)
Average of 1-5 years post-move	0.391	0.121	0.684
	(0.049)	(0.025)	(0.107)
Enrollees	76,728	45,767	30,907
Enrollee-years	586,248	342,501	243,288

Notes: Table reports the average of coefficients and their bootstrapped standard errors (in parentheses) between one to five years pre-move and one to five years post-move as estimated from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—where the weights on the regression are given by the number of movers observed in each cohort. Estimates are shown for the baseline sample, naive enrollees, and prior users observed in the year before move for whom we are able to observe matched non-movers. Each panel shows a single deviation from our baseline specification (baseline shown in Panel A). Panel B restricts to movers with 12 months of Part D coverage for each year they are observed in the sample, and Panel C restricts to all movers who did not die from 2006 to 2019. Opioid naive enrollees are those with no opioid use in relative year -1, while prior users filled at least one opioid prescription in that year. Standard errors are computed using 50 iterations of the Bayesian bootstrap.

Table A.15: Commuting Zone Estimation Sample Summary Statistics

	Movers	Non-Movers
Sample Used to Construct Moments		
Number of CZs in Sample	91	
Number of States in Sample	36	
Number of Enrollees	46,241	1,585,725
Number of Enrollee-Years	336,624	11,440,375
Number of Cohorts	21,811	–
Average Risky Use Rate	3.6%	3.4%
Moments Targeted in Model Estimation		
Number of Moments	172,259	1,274
Median Number of Enrollee-Years per Moment	1	6,554
Mean Number of Enrollee-Years per Moment	2.0	8,980
Minimum Number of Enrollee-Years per Moment	1	1,341
Maximum Number of Enrollee-Years per Moment	102	70,700

Notes: Table presents summary statistics on the empirical moments that are targeted in the estimation of the model for a subset of commuting zones as well as the sample used to construct those moments. More details on the 91 commuting zones in this sample are provided in Appendix C. Mover enrollee-years observed during the year of move are excluded from the moments used in estimation. The empirical moments are average risky use rates constructed at the cohort – relative-year level, where cohorts are defined by a mover's origin, destination, and move year.

Table A.16: Medicare Elderly Sample Summary Statistics

	(1) Movers	(2) Non-movers
Female	63%	60%
White	85%	81%
Black	8%	9%
Medicaid	17%	19%
Age:		
Average age	76.5	75.5
Region:		
Northeast	16%	19%
West	26%	23%
Midwest	17%	23%
South	41%	35%
Opioid use:		
Any opioids	22.2%	20.4%
Prescriptions in year before move ("prior user")	23.4%	
No prescriptions in year before move ("opioid naive")	76.6%	
Risky use	0.4%	0.4%
Number of enrollee-years	2,062,303	65,143,124
Number of enrollees	291,510	10,166,554

Notes: Table presents summary statistics for an alternative sample of Medicare enrollee-years in which the enrollee was 65+. All rows except for the number of enrollee-years and enrollees report the share of enrollees or enrollee-years within the given population with the indicated characteristic. "Any Opioids" and "Risky Use" are averaged over all enrollee-years, while all other statistics are averaged at the enrollee-level, with "Region," "Medicaid," and "Age" measured in a reference year. This reference year is the year before move for movers, and a randomly assigned year for non-movers such that the distribution of reference years for non-movers mirrors that of movers. More details about the construction of this sample are provided in Appendix C.

Table A.17: Medicare Elderly Event Study Coefficients - Rates of Risky Use

	(1)	(2)	(3)
	All	Naive	Prior User
1 year post-move	0.282	0.037	0.889
	(0.065)	(0.020)	(0.233)
5 years post-move	0.597	0.094	1.914
	(0.153)	(0.047)	(0.524)
Enrollees	291,510	223,877	67,612
Enrollee-years	2,062,303	1,579,471	482,674
Average Risky Use Rate 1 year pre-move	0.005	0.000	0.023

Notes: Table reports the coefficients and their bootstrapped standard errors (in parentheses) one and five years post-move as estimated from equation (3) of μ_r—the coefficient on year relative to move interacted with $\hat{\delta}_{cr}$—on the Medicare elderly sample described in Appendix C. The weights on the regression are given by the number of movers observed in each cohort. Estimates are shown for the baseline sample ("All"), naive enrollees, and prior users, for movers observed in the year before move for whom we are able to observe matched non-movers. Opioid naive enrollees are those with no opioid use in relative year -1, while prior users filled at least one opioid prescription in that year. Standard errors are computed using 50 iterations of the Bayesian bootstrap.

Table A.18: Medicare Elderly Estimation Sample Summary Statistics

	Movers	Non-Movers
Sample Used to Construct Moments		
Number of Enrollees	291,510	10,166,554
Number of Enrollee-Years	1,810,352	65,143,124
Number of Cohorts	21,225	–
Average Risky Use Rate	0.4%	0.4%
Moments Targeted in Model Estimation		
Number of Moments	163,280	714
Median Number of Enrollee-Years per Moment	4	55,951
Mean Number of Enrollee-Years per Moment	11.1	91,236.9
Minimum Number of Enrollee-Years per Moment	1	2,707
Maximum Number of Enrollee-Years per Moment	967	805,098

Notes: Table presents summary statistics on the empirical moments that are targeted in the estimation of the model for our estimation on the Medicare elderly sample, as well as the sample used to construct those moments. The sample used to construct the moments is simply the sample of all movers presented in Table A.16 column 1, excluding enrollee-years observed during the year of move. The empirical moments are average risky use rates constructed at the cohort – relative-year level, where cohorts are defined by a mover's origin, destination, and move year. More details about the Medicare elderly sample are presented in Appendix C.

Table A.19: Monte Carlo Estimations: The Effect of Counterfactuals on Risky Use

	Predicted reduction in rate of risky use from a one standard deviation reduction in...	2006 - 2019 *(True Value Below)*
(1)	Person-based Parameters (η_i)	36.8% *(32.6%)*
(2)	Place-based Parameters (γ_j, π_j)	52.5% *(41.3%)*
	(a) Availability (γ_j)	19.2% *(13.9%)*
	(b) Place-based Addiction Transitions (π_j)	41.2% *(31.9%)*
(3)	All Addiction Transitions (π_j, η_i)	61.2% *(51.0%)*

Notes: Table reports the percent reduction in predicted overall risky use rates from 2006 to 2019 from a one standard deviation reduction in parameters. Each counterfactual policy is to reduce parameters in the relevant channel(s) by a standard deviation, except for the probabilities of transitioning out of addiction, which we increase by a standard deviation. The overall risky use rate is computed by taking a weighted average of predicted risky use over all state-years in our non-mover sample, weighting by the number of non-movers observed in each state-year. 95% confidence intervals are computed using 50 iterations of the Bayesian bootstrap. Throughout, we maintain our standard parameter bounds of [0,1] by taking the minimum (maximum) of the reduced (increased) parameter and the bound as necessary.

Table A.20: Monte Carlo Estimations: Yearly Effects of Counterfactual Policies Targeting Risky Use

	Targeting Place-Based Addiction Channel *(True Value Below)*	Targeting Place-Based Availability Channel *(True Value Below)*	Targeting Both Place-Based Channels *(True Value Below)*
2006	0.0%	28.4%	28.4%
	(0.0%)	*(19.0%)*	*(19.0%)*
2007	15.6%	24.9%	36.7%
	(10.5%)	*(17.6%)*	*(26.2%)*
2008	25.8%	19.4%	40.2%
	(18.3%)	*(13.7%)*	*(29.4%)*
2009	32.9%	16.8%	44.2%
	(24.1%)	*(11.9%)*	*(33.2%)*
2010	38.0%	14.9%	47.3%
	(28.6%)	*(10.6%)*	*(36.2%)*
2011	41.9%	13.9%	50.0%
	(32.2%)	*(10.0%)*	*(39.0%)*
2012	44.9%	14.0%	52.6%
	(35.1%)	*(10.1%)*	*(41.6%)*
2013	47.4%	14.7%	55.2%
	(37.4%)	*(10.6%)*	*(44.0%)*
2014	49.5%	16.0%	57.6%
	(39.3%)	*(11.6%)*	*(46.3%)*
2015	51.3%	17.5%	59.8%
	(40.9%)	*(12.7%)*	*(48.4%)*
2016	52.8%	20.4%	62.4%
	(42.3%)	*(14.9%)*	*(50.8%)*
2017	54.0%	25.7%	65.8%
	(43.3%)	*(19.0%)*	*(54.1%)*
2018	55.0%	32.9%	69.7%
	(44.2%)	*(24.5%)*	*(57.8%)*
2019	55.8%	43.5%	74.8%
	(44.9%)	*(32.8%)*	*(62.7%)*

Notes: Table reports the estimators of the yearly effect of counterfactuals constructed from averaging one hundred Monte Carlo simulations. The true effect—which is identical to the estimates from Table A.6 because the same underlying parameters were used as the baseline for the Monte Carlo—is presented in parentheses under each estimate for comparison. Estimates are presented as the yearly percentage reduction in risky use rates driven by the effect of counterfactual policies that lower place-based addiction and place-based availability parameters as described in Section 6.3 and presented in Figure 10. Each counterfactual policy is to reduce parameters in the relevant channel(s) by a standard deviation except for the probabilities of transitioning out of addiction, which we increase by a standard deviation. The yearly risky use rate is computed by taking a weighted average of predicted risky use over all states in a given year in our non-mover sample, weighting by the number of non-movers observed in each state. Throughout, we maintain our standard parameter bounds of [0,1] by taking the minimum (maximum) of the reduced (increased) parameter and the bound as necessary.