What Drives Prescription Opioid Abuse?
Evidence from Migration

Amy Finkelstein
MIT and NBER

Matthew Gentzkow
Stanford and NBER

Dean Li
MIT

Heidi Williams
*Stanford and NBER**

July 2021

Abstract

We investigate the role of person- and place-specific factors in the opioid epidemic by developing and estimating a dynamic model of prescription opioid abuse. We estimate the model using the relationship between cross-state migration and prescription opioid abuse among adults receiving federal disability insurance from 2006 to 2015. Event studies suggest that moving to a state with a 3.5 percentage point higher rate of opioid abuse (roughly the difference between the 20th and 80th percentile states) increases the probability of abuse by 1 percentage point on-impact, followed by an additional increase of 0.30 percentage points per subsequent year. Model estimates imply large place effects in both the likelihood of transitioning to addiction and the availability of prescription opioids to the addicted. Equalizing place-based factors would have reduced the geographic variation in opioid abuse by about 50 percent over our 10-year study period. Reducing place effects on addiction transitions to the 25th percentile would have twice the impact on opioid abuse after 10 years as the analogous reduction in place effects on availability to addicts, though the comparison is reversed in the first few years.

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

*We thank the National Institute on Aging (Finkelstein R01-AG032449), the National Science Foundation (Williams, 1151497), the Social Security Administration (Finkelstein RDR18000003) and the Stanford Institute for Economic Policy Research (SIEPR) for financial support. We are grateful to Alden Cheng, Paul Friedrich, Ken Jung, Geoffrey Kocks, Abby Ostriker, Anna Russo and Yuci Zhou for excellent research assistance, and to Jonathan Gruber, Ellen Meara and Molly Schnell for helpful discussions. 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. 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 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 homicide deaths, and an order of magnitude higher than cocaine-related deaths at the height of the 1980s crack epidemic (GAO 1991; Kochanek et al. 2019; NCHS 2021).

A distinguishing feature of the epidemic is the role played by legal prescription opioids such as oxycodone (Oxycontin) and hydrocodone (Vicodin). The opioid crisis emerged from physicians' increased willingness to prescribe legal opioids, which in turn was driven by pharmaceutical firms' marketing efforts and the medical profession's decision to recognize pain as a "fifth vital sign" (Quinones 2016; Case and Deaton 2020; Maclean et al. 2020). Prescription opioids continue to account for a large share of opioid deaths today (CDC 2017a).

In this paper, we estimate a dynamic model of prescription opioid abuse that allows us to unpack the drivers of the opioid epidemic along two key 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 abuse, 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 et al. 2017; Schnell and Currie 2018; Eichmeyer and Zhang 2021), the availability of unscrupulous providers and "pill mill" pain clinics (Lyapustina et al. 2016; Rutkow et al. 2015), and the presence of policies such as triplicate prescription programs (Alpert et al. 2019).[1]

Second, we distinguish between place factors that impact the likelihood that individuals transition to and from addiction and place factors that influence the ease with which existing addicts can obtain prescriptions. The former, which we refer to as the *addiction channel*, would include factors like the willingness of physicians to give first-time prescriptions to patients suffering from pain, and it would be particularly influenced by policies like prescribing limits for such opioid "naive" patients (Sacks et al. 2021). The latter, which we refer to as the *availability channel*, would include factors like the presence of pill mills, and it would be particularly affected by policies such as prescription monitoring programs (Buchmueller and Carey 2018;

[1] Other examples of policies include naloxone access laws (Doleac and Mukherjee 2018; Rees et al. 2019), opioid prescribing limits (Sacks et al. 2021), and prescription monitoring programs intended to limit abuse (Kilby 2015; Meara et al. 2016; Buchmueller and Carey 2018; Kaestner and Ziedan 2019). Area-specific peer effects (Khan et al. 2019; Powell et al. 2020) will also be captured as place factors in our framework.

Kaestner and Ziedan 2019; Meara et al. 2016; Kilby 2015) and pill mill laws (Lyapustina et al. 2016; Rutkow et al. 2015).

Decomposing the role of these different factors is important for understanding how different public policies could have affected the epidemic. It is also essential for understanding the wide geographic variation that has been observed in the intensity of the epidemic. Prescribing per capita varies by a factor of more than two between the 10th and 90th percentile states (McDonald et al. 2012; Schieber et al. 2019), and the crisis has been especially severe in Appalachia, New England, and the West (Katz 2017; Guy et al. 2017).

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 an individual abuses opioids depends on her state of addiction as well as place-specific supply (i.e., availability) factors. The model yields rich predictions for variation in the dynamics of the epidemic across space and time. It suggests that changes in abuse rates when people move can identify the relative importance of person and place-specific factors, with the evolution of abuse rates following moves further distinguishing between the place-based addiction and availability channels.

We estimate the model using data on prescription opioid use from 2006 to 2015 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 receive an opioid prescription each year (Meara et al. 2016; Morden et al. 2014), a rate considerably higher than the rest of the population.[2] We take advantage of rich panel data on prescription drug use for about 2.4 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 in driving changes in behavior around the times when people move. At the same time, it may also attenuate some of the variation in economic circumstances highlighted by Case and Deaton (2015; 2017; 2020).

To define prescription opioid abuse, we follow a standard approach in the literature (e.g., Ciesielski et al. 2017; Meara et al. 2016; Larrick 2014) and measure abuse for each individual-year with an indicator for prescriptions that result in a daily morphine-equivalent dose (MED) of more than 120 mg in any calendar quarter. This measure is strongly correlated with adverse outcomes such as opioid poisonings and opioid dependence in both our data and in other populations.

[2]The share of opioid prescriptions accounted for by SSDI recipients is almost three times higher than their population share. SSDI recipients account for less than 5 percent of the adult population, and we estimate that adult SSDI recipients account for about 13 percent of all opioid prescriptions (Autor 2015; SSA 2017; CDC 2017b; authors' calculations). We calculate SSDI recipient opioid use by multiplying the average opioid use per SSDI recipient on Medicare Part D (observed in our data) by the total number of SSDI recipients.

We begin with event study estimates of changes in opioid abuse when individuals move across state lines. We find that individuals moving to higher-opioid abuse areas immediately begin abusing opioids at higher rates, and those moving to lower-abuse areas immediately begin abusing at lower rates. These impacts increase in the subsequent years post-move. For example, we estimate that moving from a state at the 20th percentile of opioid abuse to a state at the 80th percentile of opioid abuse (about a 3.5 percentage point increase) causes an immediate jump in opioid abuse of approximately 1 percentage point followed by a continued increase of about 0.30 percentage points per subsequent year.

Viewed through the lens of our model, the large on-impact effects of moving are consistent with a large role for place-specific factors that affect availability. The continued post-move convergence in abuse rates additionally suggests an important role for place-specific factors that affect transitions to addiction. Consistent with this interpretation, when we proxy for an individual's state of addiction with her prior opioid use, the results suggest that for prior opioid users, moving to a higher-abuse state is associated with a sharp on-impact increase in prescription opioid abuse along with further increases post-move, whereas for opioid-naive patients, there is little or no on-impact change, but a steady increase post-move.

We estimate the model using method of moments. We use moments that measure the average rate of opioid abuse each year before a move and each year after a move, separately for about 12,500 mover cohorts, defined by their origin and destination states and the year in which their move occurs. Event studies simulated based on the estimated parameters match the observed event studies reasonably well.

The estimated parameters allow us to analyze the impact of different channels in contributing to prescription opioid abuse. This reveals a quantitatively important role for both person-specific and place-specific factors, and among place-specific factors, a particularly large role for the addiction channel. We estimate—counterfactually—that equalizing place-specific factors during our study period would have eliminated about half of the cross-state variation in opioid abuse. We also estimate the equalizing addiction transition probabilities—which include both place and person components—would have reduced the cross-state variation in opioid abuse by three times more than eliminating place-specific availability of opioids to addicts.

Finally, we consider the relative impact of policies that reduce the place-specific availability channel—for example, through pill mill laws and prescription monitoring programs—to that of policies that reduce the place-specific addiction channel—for example through policies that regulate physician prescribing behavior for opioid naives. Specifically, we examine the impact on opioid abuse of reducing states' place-based addiction transition rates to the 25th percentile of the state distribution and of reducing states' place-based availability to addicts to the 25th percentile of the availability distribution. We find that by the

end of our 10-year study period, the reduction in the place-based addiction channel would have reduced opioid abuse by twice the amount as a reduction in the availability channel. At the same time, our estimates reveal an important temporal dimension to this trade-off, with the impact of reducing the availability channel relatively greater in the short run and the impact of reducing the addiction channel relatively greater in the longer-run.

An important limitation to these conclusions is that the place-specific factors that we measure only capture changes in prescription opioid abuse; our data do not permit a direct examination of potential substitution to illegal forms of opioids such as heroin. Further analysis of the determinants of illegal opioid use is an important area for future work, particularly as illegal opioids have started to play a larger role in the epidemic, especially in more recent years beyond our study period (Maclean et al. 2020). A second important limitation is that our focus on the SSDI population means that variation in income will be smaller in our setting than in the general population, limiting our ability to pick up the role of both person and place-specific drivers of individual economic circumstances.

Our results speak to an ongoing debate over the relative importance of different factors in driving the opioid epidemic. Much of this work has been based on examination of the correlations between patterns of rising opioid abuse and changes in other factors. Case and Deaton (2015; 2017; 2020) and others have pointed to the important role of individuals' economic circumstances and related demand-side factors (Carpenter et al. 2017; Hollingsworth et al. 2017; Dean and Kimmel 2019; Venkataramani et al. 2020). Others have argued that the patterns of increased drug deaths and opioid abuse across states suggest an important role for the availability of drugs and little, if any, role for economic decline (Ruhm 2018; Currie et al. 2019; Currie and Schwandt 2020; Cutler and Glaeser forthcoming).[3] Charles et al. (2019) conclude that both economic conditions and opioid supply have played in rising local opioid use and deaths.

We develop an explicit model of opioid addiction, and we use a different empirical strategy, based on individuals moving across states, to separate person-specific and place-specific factors, and to distinguish among different types of place-specific factors. Our strategy is similar to one we have used in prior work to understand the determinants of geographic variation in the elderly's health care utilization (Finkelstein et al. 2016), but extends the approach to allow for non-separability in the outcome equation of person-specific and place-specific factors, as well as allowing for distinct place-specific channels.[4] This mover-based empirical

[3] Other work suggests that economic decline may be a result of opioid abuse itself (Laird and Nielsen 2016; Krueger 2017; Aliprantis et al. 2019; Harris et al. 2020).

[4] Finkelstein et al. (2016) in turn draw on earlier work in the health care space that exploited patient migration (Song et al. 2010) and physician migration (Molitor 2018). 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.

approach fits more broadly into a large and growing prior 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 economic outcomes (e.g. Card et al. 2013; Chetty et al. 2014; Chetty and Hendren 2018a,b). Most of this prior literature has assumed separability between person-specific and place-specific factors in the outcome equation.[5]

Finally, our analysis of the sources of geographic variation in prescription opioid abuse contributes 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 important role we estimate for place-specific factors that affect transitions to addiction is consistent with the existing widespread evidence of substantial geographic variation in treatment practices that involve tradeoffs across different objectives, such as whether to order a test or whether to engage in aggressive treatment or "watchful waiting" in response to early detection of prostate cancer.

The rest of the paper proceeds as follows. Section 2 describes the setting and data. 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 estimates the model and presents our analysis of the impacts of counterfactual place-based policies. The last section concludes.

2 Setting and Data

We study prescription opioid abuse over the ten year period from 2006 to 2015. During this time period, the opioid epidemic was already firmly established, but the crisis was still primarily one of *prescription* abuse (see Appendix Figure A.1).

We analyze prescription opioid use among 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. Over 8 million disabled workers (and 2 million spouses and dependents) received SSDI benefits in 2010, and SSDI expenditures comprised over seven percent of federal non-defense spending. About one-third of SSDI expenditures reflect Medicare costs (Autor 2015). Once on SSDI, it is rare for individuals to exit SSDI and return to the labor force; most exits are either due to death or to reaching age 65 and so qualifying for Social Security old-age benefits (Autor and Duggan 2006). While on SSDI, recipients' income and labor market status tends to be stable; they are

[5]The study of habit formation in consumer purchases by Bronnenberg et al. (2012) is one exception.

not allowed to have monthly labor market earnings above a low limit, which was $1,090 per month in 2015 (SSA 2017).

Several features of this population make it particularly attractive for our analysis. First, the stringent limits on labor market activity allow us to abstract from potential changes in opioid abuse driven by large changes in individual employment or income around the time they move. Second, rates of opioid use are high among this population (Meara et al. 2016; Morden et al. 2014). And third, through the Medicare data we are able to observe detailed, panel data on their prescription opioid use.

We use data from SSDI recipients included in a 20 percent random sample of Medicare beneficiaries.[6] We focus on the approximately three-quarters of SSDI recipients who are also enrolled in Medicare Part D. This voluntary, heavily subsidized prescription drug benefit program has been available to Medicare enrollees since 2006. They can enroll in it either 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. The Part D enrollment rate among SSDI Medicare recipients is slightly higher than the 60 percent rate among elderly Medicare recipients (Cubanski et al. 2016). Adults with sufficiently low earnings prior to the onset of disability also quality for Medicaid - the public health insurance program for qualifying low income adults; dual eligibility for Medicaid and Medicare ensures that beneficiaries have Part D coverage, and face no out of pocket expenses for covered health care.

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

Sample construction We define all outcome variables for year t to be aggregates of claims from January through December of year t, and analyze data for 10 years: $t = 2006$ through $t = 2015$. From our original

[6]Unlike Medicare for the elderly, enrollment in SSDI is not limited to those over 65 years of age, so our data includes a wide range of ages.

[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 2009-2015. 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. More information is available here: `https://www.resdac.org/articles/medicare-eligibility-and-enrollment-files-rif-versions`.

20 percent sample of all Medicare enrollees (15.6 million enrollees; 103 million enrollee-years) we limit ourselves to the 3.5 million enrollees (23.5 million enrollee-years) whose original source of Medicare eligibility was from SSDI. We further restrict to enrollee-years with 12 months of Medicare Part D coverage, which excludes approximately 930,000 enrollees and 8.8 million enrollee-years.

Our baseline geographic unit of analysis is a state, although we show robustness of our descriptive analysis to other levels of geography such as county and 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 215,000 enrollees who change their state of residence at least once, we impose several additional restrictions to arrive at our sample of 90,000 movers. We exclude the approximately 25,000 movers who moved during the last year in the sample since we cannot observe their post-move behavior. We exclude about 65,000 movers who changed their state of residence more than once between April 2006 and December 2015. We limit analysis to the five years pre- and the five years post-move, thus excluding about 90,000 enrollee-years outside of this window. Finally, we follow the approach developed in Finkelstein et al. (2016) and exclude about one-quarter of the remaining movers whose share of prescription claims in their destination state, among prescription claims 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] 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.

Our final sample contains 90,890 movers (521,523 enrollee-years) and 2,325,094 non-movers (13,349,773 enrollee-years).

Measuring opioid abuse 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 opioid abuse when a patient may be feigning pain (Parente et al. 2004).

[8]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 et al. (2016) for more discussion.

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

Our primary measure of abuse 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 quarter of the year. Our use of this measure follows Ciesielski et al. (2017), Larrick (2014), Meara et al. (2016), and Sullivan et al. (2010). The dosage required for abuse is above the 96th percentile of person-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.

We also show robustness of our descriptive results to two other measures 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. 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. Our primary abuse measure is the most predictive of subsequent adverse outcomes in our data.

Finally, for some of our descriptive analyses, we will find it useful to distinguish among patients based on their prior opioid history. Following the standard definitions in the literature (Deyo et al. 2017; Sun et al. 2016), we define an "opioid naive" as a migrant who filled no opioid prescriptions in the year before the move ($r = -1$), and a "prior user" as a migrant who filled an opioid prescription in the year before the move.

Summary statistics We report summary statistics for our study population of movers in Table 1, column 1. Slightly over one-half of movers are female, and approximately three-quarters are white. Almost three-fifths receive Medicaid, and their average age is 56. 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 (Meara et al. 2016; Morden et al. 2014). About five percent of enrollee-years meet the

definition for abuse. Those whom we can characterize are roughly evenly divided between opioid naives and prior users; almost one-fifth of mover-years have no Part D data in the year before move and are therefore not characterized as an opioid naive or a prior user in the year prior to move. Relative to non-movers (column 2), movers are somewhat more likely to be female, white, on Medicaid, and younger. Movers also exhibit somewhat higher rates of prescription opioid use and abuse.

Among our movers, 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 1,462 movers and the mean state receives 1,779. Florida is the most common destination state (about 13% of movers) and also the most common origin state (8.7% of movers). The least common destination is the District of Columbia (0.2% of movers) and the least common origin is North Dakota (0.2% of movers).

Figure 1 shows the distribution of the rate of opioid abuse across states for our full sample. 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 abuse rate in New Hampshire is more than twice that of the neighboring state of Massachusetts, while in the Midwest, the abuse rate in Montana is more than two and a half times higher than that of the neighboring state of North Dakota.[9]

3 Model

We build a dynamic model of opioid addiction and abuse. 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. We let $r(c,t) \equiv t - m(c)$ denote the year relative to a cohort's move.

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 state evolves stochastically with 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 $\pi_j^+ + \eta_i^+$ and the probability

[9]We cannot compare the pattern of abuse as we measure it in our population to that of the general population. However, in Appendix B we have confirmed 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 opioid prescriptions is substantially higher for our disabled population.

of transitioning out of addiction is given by $\pi_j^- + \eta_i^-$, where the π and η terms represent place and person-specific factors respectively. Thus,

$$Pr(a_{i,t} = 1 | j(i,t) = j) = \begin{cases} 1 - \pi_j^- - \eta_i^- & \text{if } a_{i,t-1} = 1 \\ \pi_j^+ + \eta_i^+ & \text{if } a_{i,t-1} = 0. \end{cases}$$

We let $\pi_j = \begin{bmatrix} \pi_j^+ & \pi_j^- \end{bmatrix}$ and $\eta_i = \begin{bmatrix} \eta_i^+ & \eta_i^- \end{bmatrix}$.

Because underlying drivers of addiction and abuse have changed dramatically over time, we do not focus on the steady state of this model but rather study the evolution of addiction patterns beginning from an initial state that we estimate. The addict share $\overline{a}_{jt}$ among non-movers in location j is defined recursively as

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

where the initial addiction states a_{i0} are parameters to be estimated and $\left|\mathscr{I}_j\right|$ denotes the number of non-movers in location j.

Our main observed outcome of interest is abuse of prescription opioids. We treat this as binary, letting $y_{it} \in \{0,1\}$ be an indicator for whether person i abuses prescription opioids in year t. We assume that the probability an addict abuses opioids is a function of place and time-specific factors, including local supply-side conditions. We assume that non-addicts never abuse opioids.

Thus opioid abuse is determined by

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

Thus expected abuse conditional on the addiction state is $a_{it}\gamma_{jt}$. As a shorthand, we refer to the parameter γ_{jt} as "opioid availability" in location j and time t.

In the empirical analysis, we will observe opioid abuse y_{it} and the location of each individual, but not their underlying addiction state a_{it}. We cannot directly observe the binary addiction state a_{it}, since not all addicts are prescribed opioids in a given year. The key identification challenge will therefore be to separately pin down the opioid availability parameters γ_{jt}, the place-specific and person-specific addiction parameters π_j and η_i, and the initial addict shares $\overline{a}_{j0}$. As we will discuss, the patterns of opioid abuse around moves to and from different locations will be crucial for identification.

A main focus of the analysis will be characterizing the component of abuse that can be attributed to the

causal effect of place. For an addict ($a_{it} = 1$), the causal effect on abuse in period t of moving in period t from an original location $o(i)$ to a new location $d(i)$ is $\left(\gamma_{d(i)t} - \gamma_{o(i)t}\right)$. The causal effect for a non-addict is zero. Thus, the average causal effect of such a move in period t for a population in which share $\overline{a}_t$ is addicted is $\overline{a_t}\left(\gamma_{d(i)t} - \gamma_{o(i)t}\right)$. Over time, the share addicted $\overline{a}_t$ will also be affected by the move via place-specific addiction parameters.

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

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

where $\mathbf{h}_{ct}^0$ 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

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

where $\overline{a}_{ct}$ is the realized share addicted in cohort c at time t. Note that, as shown in equation 2, the period-t average treatment effect on movers consists of two distinct channels, reflecting both the impact on the addict share ($\overline{a}_{ct}$) and on abuse conditional on addiction. To see this more clearly, we can rewrite the term for years after the move as

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

where the first term shows the treatment effect due to the addiction channel and the second term shows the treatment effect due to the availability channel.

4 Descriptive Evidence

We present descriptive evidence on how opioid abuse changes with the timing and direction of moves, and how these patterns vary across individuals with different probabilities of addiction. We show how these patterns can be interpreted in light of the model in Section 3 to provide an initial sense of the relative importance of place- and person-based factors in contributing to geographic variation in opioid abuse, as well as the relative importance of the two different place-based channels, availability and addiction.

4.1 Event Study Specification

Our event study specification follows recent work emphasizing the importance of allowing for flexibility in treatment effects across different treated units (e.g. Sun and Abraham 2020, Callaway and Sant'Anna 2020, and de Chaisemartin and D'Haultfœuille 2020). We begin by regressing y_{it} on location-year fixed effects and five-year age bins x_{it} in the sample of non-movers (our "control" group). We then use these estimated coefficients to compute predicted values for each mover in our data, assigning them the location-year fixed effects associated with their origins. We then subtract these predictions from y_{it} to form residuals $\tilde{y}_{it}$. Using only non-movers for the prediction step guarantees that the location-year fixed effects and age coefficients are not contaminated by the treatment effects themselves.

We then estimate the following event study on the sample of movers only, separately by move year m:

$$\tilde{y}_{it} = \mu_{r(c,t)} \hat{\delta}_{ct} + \varepsilon_{it}. \tag{3}$$

The key variable of interest is $\hat{\delta}_{ct}$, which is our estimate of the "size" of a move. Specifically, for each cohort c, we define δ_{ct} to be the difference in the average rates of opioid abuse among non-movers in period t between the cohort's destination $d(c)$ and origin $o(c)$:

$$\delta_{ct} = \bar{y}_{d(c)t} - \bar{y}_{o(c)t}.$$

We calculate its sample analog $\hat{\delta}_{ct}$on the full sample of non-movers shown in Table 1, column 2.

The event study coefficients $\mu_{r(c,t)}$ are the main estimates of interest; they show the pattern of opioid abuse around moves as a function of the difference in the abuse rate of the destination relative to the origin. Differencing out predicted abuse rates by origin by year allows for differential trends in opioid abuse rates across areas, while differencing out predicted abuse rates by five-year age bin allows for trends in abuse by age. Because we estimate the event study separately by move year, we are allowing for heterogeneous treatment effects by the timing of the treatment. Following Callaway and Sant'Anna (2020), we summarize the event study coefficients that we estimate separately for each move year m with a simple weighted (by sample size) aggregation of these coefficients.

To facilitate precise interpretation of the event study, we do not include additional controls in equation (3). We show below that the results are robust to adding individual fixed effects (which would allow for fixed differences across movers in the probability of opioid abuse) as well as indicators for the year relative to the move year (which would allow for potential direct impacts of moving itself on opioid abuse).

Interpretation

Under the assumption that all movers are randomly drawn from the population of their origin location, the event study coefficients μ_r can be interpreted as a weighted average of cohorts' treatment effects T_{ct} relative to the average difference in outcomes between their origin and destination. We will examine the pattern of coefficients prior to move as one way to investigate the validity of this random selection assumption.

Proposition 1. *The probability limit of the event study coefficients $\hat{\mu}_r$ in equation* (3) *estimated on the sample of movers with move year m is*

$$\mu_r = \sum_{c \in \mathscr{C}_m} w_c \frac{T_{ct}}{\left(\overline{y}_{d(c)t} - \overline{y}_{o(c)t}\right)} \tag{4}$$

where $t = m + r$, $\mathscr{C}_m$ is the set of cohorts with move year m, $w_c \equiv \frac{|\mathscr{I}_c| \cdot \delta^2_{c(i),t}}{\sum_{c \in \mathscr{C}_m} |\mathscr{I}_c| \cdot \delta^2_{c(i),t}}$ is the weight given to each cohort c with $|\mathscr{I}_c|$ denoting the number of individuals in that cohort.

Proof. See Appendix C. □

The weights w_c in Proposition 1 are increasing in the number of individuals in a cohort and the difference in abuse rates between their origin and destination.

To build intuition for what we can learn from the pattern of event study coefficients, consider a population consisting of a single cohort c with origin $o(c)$ and destination $d(c)$. Recall from equation (2) that T_{ct} must be zero in years prior to the move, and thus μ_r for $r \leq 0$ must be equal to zero as well. In years following the move we have $T_{ct} = \overline{a}_{ct}\gamma_{d(c)t} - \overline{a}_{o(c)t}\gamma_{o(c)t}$ and $\overline{y}_{d(c)t} - \overline{y}_{o(c)t} = \overline{a}_{d(c)t}\gamma_{d(c)t} - \overline{a}_{o(c)t}\gamma_{o(c)t}$. Thus, we can rewrite μ_r for $r > 0$ as

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

We consider four special cases.[10]

Case 1: Only Person Effects Suppose, first, that neither availability effects γ_{jt} nor transition probabilities π_j vary across the origin and destination. Then all geographic variation is due to person effects. Abuse rates differ across places only because of differences in the distribution of initial addiction states a_{i0} and addiction propensities η_i. In this case, we have $\gamma_{d(c)t} = \gamma_{o(c)t}$ and also $\overline{a}_{ct} = \overline{a}_{o(c)t}$ (since addiction rates do not depend

[10]Note that, 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 discuss the effects of an instantaneous move. However in the visualization of our various examples in Figure 2 below, we replicate uncertainty around the mover's location in relative year 0.

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_{jt} differ across the origin and destination but the transition probabilities π_j and the distributions of a_{i0} and η_i do not. Abuse rates differ across places only because of differences in the ease with which addicts can obtain opioids, and all geographic variation is due to place effects in availability. In this case, we have $\overline{a}_{ct} = \overline{a}_{o(c)t} = \overline{a}_{d(c)t}$, since the evolution of addiction rates does not depend on location. Thus, $\mu_r = 1$ for all $r > 0$, and 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 both γ_{jt} and the distributions of person-specific factors a_{i0} and η_i vary across the origin and destination, but continue to assume the transition probabilities π_j do not vary across locations. Then we have $\overline{a}_{ct} = \overline{a}_{o(c)t}$ but $\overline{a}_{o(c)t} \neq \overline{a}_{d(c)t}$. In this case, we can write the event-study coefficient for $r > 0$ as

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

Here, the first (addiction) term in the denominator is due only to differences person-specific factors and the second (availability) 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 this effect is the same as the sign of the difference in addiction rates $\left(\overline{a}_{d(c)t} - \overline{a}_{o(c)t}\right)$, this share will lie between zero and one. The event study plot will show an on-impact jump equal to this share and have 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_j to also vary across the origin and destination. This means that $\overline{a}_{ct}$ will differ from both $\overline{a}_{o(c)t}$ and $\overline{a}_{d(c)t}$, and we cannot simplify the expression for μ_r beyond equation (5).

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

Note, first, that immediately after the move $\overline{a}_{ct}$ will be close to $\overline{a}_{o(c)t}$. This means that the on-impact jump in the event-study plot will have the a similar interpretation as in case 3, roughly interpretable as the share

of differences in outcomes that are due to place-specific availability factors. Next, consider the changes in $\overline{a}_{ct}$ post-move. If most variation in addiction rates is due to person-specific factors a_{i0} and η_i, $\overline{a}_{ct}$ will remain close to $\overline{a}_{o(c)t}$ and the event study plot will remain flat post-move as in case 3. If a large share of variation is due to the place-specific factors π_j, $\overline{a}_{ct}$ will gradually increase following the move from a level close to $\overline{a}_{o(c)t}$ in the direction of $\overline{a}_{d(c)t}$. This positive post-move trend in the event study plot is a signature of variation in π_j playing an important role. If that variation is sufficiently important that the addiction rates among movers eventually converge to the destination level $\overline{a}_{d(c)t}$, the event study coefficients μ_r will eventually converge toward 1. Panel D of Figure 2 illustrates this case.

4.2 Results

Aggregate Event Studies

Figure 3 shows the distribution of $\hat{\delta}_{ct}$ across individual movers and years. The mean value is close to zero and the distribution is roughly symmetric, implying that moves from low- to high-abuse states are as common as moves from high- to low-abuse states. About one-tenth of movers have a $\hat{\delta}_{ct}$ of 3.5 percentage points or more in absolute value, which is roughly the equivalent to the difference between the 20th and 80th percentile state-level average abuse rates over our full sample period.

Figure 4 shows our aggregate event-study results, plotting a simple weighted aggregation of the coefficients $\hat{\mu}_{r(i,t)}$ as estimated across move years and normalizing the coefficient for relative year -1 to zero. This normalization adjusts for any time-invariant migrant-specific differences in the level of the outcome that are correlated with the migrant's origin and destination states.

Figure 4 indicates little systematic trend pre-move, which is supportive of our identifying assumption that matched non-movers serve as valid controls. It also shows the two distinct features that we might expect based on our most general Case 4 above: an immediate jump in abuse upon move and a gradual post-move convergence. As discussed, the magnitude of the jump in abuse upon move provides a rough measure of the share of the difference between a typical origin and destination attributable to differences in opioid availability to addicts. This might reflect place-specific factors such as the availability of "pill mill" pain clinics or prescription monitoring programs intended to limit abuse. In addition, the significant post-trends that we see suggest that place-specific transitions to addiction play a significant role. This might reflect place-specific factors such as peer effects, social learning, and the propensity of local physicians to initiate opioid prescriptions. These features suggest that both transitions to addiction and opioid availability to addicts play a significant role in explaining cross-sectional variation.

Column 1 of Table 2 provides a quantification of these estimates, summarizing the average values of $\hat{\mu}_r$ at various time horizons r post-move. One year after moving, we estimate that this coefficient is 0.270 (standard error = 0.037). Five years after moving, this estimate grows to 0.596 (standard error = 0.029). To illustrate these dynamics, consider locations with a stable 3.5 percentage point difference between non-movers in the rate of opioid abuse—approximately equivalent to the difference between the 20th percentile state and the 80th percentile state in average abuse rates over our full sample period. A move between these two locations would be associated with an immediate increase in the likelihood of abuse of approximately 1 percentage point in the year after moving, or about one-quarter of the cross-sectional gap between these areas. On average, the likelihood of abuse would continue increasing by about 0.30 percentage points each subsequent year after the move, with approximately three-fifths of the cross-sectional gap closed by five years after the move.

Figure 5 provides an alternative way to visualize changes in the rate of opioid abuse around moves. We consider all mover observations between 1-5 years pre-move and 1-5 years post-move. We begin by dividing these movers into ventiles according to the average destination-origin differences in non-mover abuse rates ($\bar{\bar{\delta}}_{ct}$) over all observed periods. We then plot the average change in opioid abuse rates over the 1-5 years post move, compared to the 1-5 years prior to move for each ventile of destination-origin differences. The results show a clear positive relationship, with the slope (0.25) similar to the size of the discrete jump upon move observed in the event study in Figure 4. 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 abuse rates, consistent with the functional form imposed by equation (3).

Prior Users and Opioid Naives

The model described in Section 3 predicts that changes in availability would directly affect addicts and be immediate upon move. Non-addicts, on the other hand, would be unaffected by the immediate changes in availability. If there were differences in place-based transitions to and from addiction, we would expect to see post-trends for both addicts and non-addicts.

While we cannot directly observe addiction states, we explore these predictions empirically by proxying for addiction with the enrollee's prior opioid history and examining how changes in opioid abuse around moves differ based on the enrollee's prior opioid history. 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 et al. 2017; Jeffery et al. 2018; Brat et al. 2018). Figure 6 shows the results from estimating equation (3) separately

for movers with pre-move opioid utilization ("prior users") and movers without prior opioid use ("opioid naives"). These results are summarized in Table (2) (columns 2 and 3) and are consistent with the intuitions from the model. Prior users show a larger and immediate jump upon move while there is no discrete change upon move for opioid naives; in the subsequent post-move years, abuse rates for both prior users and opioid naives increase gradually.

Total opioid abuse

Our estimates measure the impact of place-specific factors on prescription opioid abuse, but do not directly speak to their impact on use of illegal opioids such as heroin. While prescription opioids were the main driver of the crisis in its early years, illegal opioid use has become increasingly important, particularly since the end of our sample (Rummans et al. 2018; Kochanek et al. 2019).

Our data do not permit us to directly examine illegal opioid abuse or total opioid abuse. However, two pieces of suggestive evidence are consistent with total opioid abuse moving in the same direction as prescription opioid abuse. First, given the highly addictive nature of illegal opioids such as heroin, we might expect significant substitution toward them when supply conditions for prescription opioids are tightened, but little or no substitution away from them when these conditions are relaxed. This would suggest that changes on moves to higher abuse states should more accurately capture changes in total abuse. Figure 5 suggests that the effects of moves to higher abuse states are roughly comparable to the effects of moves to lower abuse states. Second, if we look at opioid poisoning events as an outcome, the results are imprecise but consistent with effects in the same direction as our main results, particularly for prior users (Appendix B).

Robustness

These descriptive results on the changes in prescription opioid abuse around moves—both in aggregate and separately based on prior opioid use—are robust to several alternative specifications (Appendix B). These include adding controls for individual fixed effects, adding controls for a series of indicator variables for the year relative to move, or adding both sets of controls. Our results are also robust to aggregating coefficients across move years that were estimated on panels balanced in event time, which rules out the possibility that our estimates are substantively affected by the changing composition of move-year cohorts and units across relative years. And they are robust to defining movers over different geographic units than states, to alternative ways of defining our sample to handle enrollee exit or entry, and to alternative approaches to defining opioid abuse.

5 Model Estimates

5.1 Estimation

Using the approximately 12,500 cohorts of movers that we observe—defined by unique combinations of origin states, destination states, and move years—we define the average rate of opioid abuse $\bar{y}_{cr}$ by cohort for relative years $r(c,t) \in [-5,-1]$ and $r(c,t) \in [1,5]$ such that:

$$Pr(y_{it} = 1 \,|\, i \in \mathscr{I}_c, r(i,t) = r) = \bar{y}_{cr}.$$

We omit moments from the move year ($r = 0$), since the enrollee is neither fully in her origin or destination in that year (see Appendix Figure A.2). We observe each cohort for an average of 6.5 years to generate approximately 80,000 moments. We denote their sample analogs in the data by $\hat{\bar{y}}_{cr}$.

We make three key assumptions in the estimation of our model. First, we assume that availability effects can be written as an additively separable function of location and time: $\gamma_{jt} = \gamma_j + \tau_t$. This ensures that geographic differences in opioid availability remain constant throughout out sample while allowing for shifting national attitudes and guidelines to affect availability across all states.

Second, we assume that movers are randomly selected among the non-movers in their origin, and so do not differ from non-movers in their origin locations in either their initial addiction states or their person-specific addiction transition factors. This is a strong assumption, but one that is supported by the lack of pre-trends in our event studies and the ability of the model to replicate the out-of-sample moments for non-movers which we demonstrate below.

Finally, for computational ease, we assume that individual-specific addiction propensities are constant within origin locations, so that $\eta_i = \eta_j$ for all i such that $o(i) = j$. This assumption allows us to derive an explicit formula to describe the evolution of average addict shares.

These assumptions imply that the addict share among movers in cohort c evolves according to

$$\bar{a}_{ct} = \begin{cases} \bar{a}_{o(c)t} & \text{if } r(i,t) < 0 \\ \bar{a}^*_c - (\bar{a}^*_c - \bar{a}_{o(c)m(c)}) \cdot (1 - v_c)^t & \text{if } r(i,t) > 0, \end{cases}$$

where $\bar{a}^*_c = \frac{\pi^+_{d(c)} + \eta^+_{o(c)}}{\pi^+_{d(c)} + \pi^-_{d(c)} + \eta^+_{o(c)} + \eta^-_{o(c)}}$ describes the steady-state addict share to which the cohort converges and $v_c = \pi^+_{d(c)} + \pi^-_{d(c)} + \eta^+_{o(c)} + \eta^-_{o(c)}$ denotes the speed at which this convergence occurs. Prior to the move, the movers' addict shares match those of non-movers in their origin. After the move, their addict shares converge toward a steady state level determined by the place-based addiction parameters $\pi_{d(c)}$ in their destination and

their own person-specific addiction propensities $\eta_i = \eta_{o(c)}$.

Availability, on the other hand, changes discretely and only once upon move. Opioid abuse among addicts is therefore given by

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

Combining these two, we can calculate average opioid abuse for cohort c in time t as

$$\bar{y}_{j(c,t)t} = \bar{a}_{ct} \cdot \gamma_{j(c,t)t},$$

where $j(c,t)$ is the location of cohort c in year t.

We estimate the parameters of this model using a three-step nonlinear weighted least squares estimator. In the first step, we begin with a set of parameter guesses. In the second step, using these parameter guesses as initial values, we minimize the weighted squared distance between our sample moments and the moments implied by our model parameters, where each moment is weighted by the observed number of movers in the data. Finally, we select the globally optimal solution. We discuss our model estimation further in Appendix D.

Identification The discussion of the event study in Section 4.1 above provides intuition for identification. Cohorts from a given origin all have identical person-specific factors but encounter different place-specific factors according to the destination to which they move. Place-specific availability effects γ_{jt} will be informed by the sharp changes in abuse rates that occur on move. Place-specific addiction effects π_j will be related to the evolution of these abuse rates post-move. Person-specific effects will be determined by patterns in abuse rates that are common to movers from a given origin regardless of their destination.

5.2 Fit

To evaluate model fit, we simulate individual-level opioid abuse for both movers and non-movers according to our estimated parameters. Following our assumption that there is no selection into who moves, we simulate an initial addiction state for movers such that $P(a_{i0} = 1 \mid o(i) = j) = \hat{\bar{a}}_{j0}$ and non-movers such that $P(a_{i0} = 1 \mid i \in \mathscr{I}_j) = \hat{\bar{a}}_{j0}$, and these initial addict shares continue to evolve according to the same estimated probabilities $\hat{\eta}_i$, and $\hat{\pi}_j$. Non-movers face the same place-specific components of addiction for all years, and addicts among non-movers abuse opioids with probability $\hat{\gamma}_j + \hat{\tau}_t$. For movers, we simulate

addiction transitions using $\hat{\pi}_{o(c)}$ as the place-specific component before the move and $\hat{\pi}_{d(c)}$ as the place-specific component after the move. Addicts among movers abuse with probability $\hat{\gamma}_{j(c,t)} + \hat{\tau}_t$. Using this simulation procedure, we generate a corresponding abuse panel for both movers and non-movers in our Medicare data and compare aggregate moments in our model simulation to the data.

The simulated moments align closely with the empirical moments (Appendix Figure A.8). This is true for both movers—on whom the model was estimated—and non-movers, who were not used in estimation and therefore provide an "out of sample" examination of fit. The time series patterns of simulated average opioid abuse rates for non-movers in the 20th, 50th, and 80th percentile states also align well with the predictions from the model (Appendix Figure A.9). The close fit of our model for non-movers is supportive of our assumption that movers are randomly selected among the non-movers in their origin.

Finally, we also compare the event study regression estimated on our simulated outcomes with the one estimated from the data in Figure 7. The model successfully matches the qualitative features of the data, and in particular, our model parameters are able to roughly replicate both the magnitude of the jump in the rate of opioid abuse upon move and the post-trends afterwards, though the observed path post move is more convex than the model would predict. Our model therefore suggests aggregate levels of importance for person- and place-specific factors in the addiction and availability channels that are broadly consistent with our reduced form results.

5.3 Results

Estimation of the model allows us to analyze the effects of various counterfactual location-based policies. In particular, we assess the role of different factors in geographic variation in abuse rates, as well as the share of abuse that could have been reduced by policies targeting different place-based factors.

These counterfactuals examine the effects of equalizing or reducing structural parameters to a given percentile in their geographic distribution. We therefore begin by defining the computation of these percentiles. For the initial probability of addiction (a_{i0}), we directly compute the median value among states. For availability effects, our assumption that these parameters are additively separable functions of location and time ($\gamma_{jt} = \gamma_j + \tau_t$) allows us to set the geographic components (γ_j) equal to their median value among states. This is equivalent to simply setting the yearly availability parameter (γ_{jt}) equal to its median value among states in each year.

Finally, following Section 5.1, we define the median values for our addiction parameters as the median steady state and convergence speed implied by the combination of addiction parameters. The steady state addict share towards which the average non-mover addict share will converge is given by $\frac{\pi_j^+ + \eta_j^+}{\pi_j^+ + \pi_j^- + \eta_j^+ + \eta_j^-}$

and the convergence speed is defined as $\pi_j^+ + \pi_j^- + \eta_j^+ + \eta_j^-$, where $\eta_i = \eta_j \; \forall \; i \in \mathscr{I}_j$. For each subset of addiction transition parameters—place-based addiction transitions (π_j) and person-based addiction transitions (η_i)—we define the steady state and convergence speed holding the other subset of addiction transition parameters constant.

Having defined this element of our counterfactuals, we begin by examining the share of the geographic variation in average abuse rates that would have been reduced over our study period if various parameters had been equalized to their median values across states in 2006 (Table 3). We find that about half (47 percent) of the cross-state variance in average opioid abuse rates over our ten-year period would have been eliminated if the place-specific effects of addiction (π_j) and yearly availability (γ_{jt}) were made equal across areas. Conversely, we find that about 30 percent of the variation would have been eliminated if the person-specific effects (addiction η_i and initial addiction state a_{i0}) were equalized across areas.[11] Differences across areas in initial addiction states (a_{i0}) may in turn reflect both differences in population characteristics and also prior influences of place.

These results are consistent with our event-study analysis, which suggested a sizable role for both person-specific and place-specific factors. When we further consider the addiction and availability channels in our model, we find that 63 percent of the state-level variation in average abuse rates over our ten-year period would have been eliminated if all parameters affecting addiction (both person-specific η_i and place-specific π_j) were equalized, whereas 14 percent of this variation would have been eliminated by equalizing the initial addict shares in each state and 26 percent would have been eliminated by equalizing only the place-specific yearly availability parameters.

These average effects over our 10-year study period mask substantial variation in the time path of the impacts of different channels. To illustrate this, Figure 8 shows the time path of the geographic variation in opioid abuse under the counterfactual equalization of different parameters. It shows, for example, that the largest driver of geographic variation in opioid abuse in the short-term—the share of addicts in the first year of our model (a_{i0})—ends up have the smallest effect on long-run geographic variation in abuse (panel C). On the other hand, equalizing addiction transition probabilities (π_j and η_i)—which has the largest effect over our sample period—does not reduce variation in the first two years (panel D).

To more concretely explore the tradeoff between short-term and long-term effects, we consider the effectiveness of policies that target different place-based channels. In particular, we compare the effectiveness of reducing yearly opioid availability parameters in states where this availability is high to corresponding re-

[11] Note that these shares need not sum to 1, both because of the non-zero correlation between person and place effects and because of their non-linear translation into abuse.

ductions in place-based addiction parameters. These two targets mirror much of recent and historical policy aimed at curbing opioid abuse. Many initial policy efforts focused on policies such as prescription monitoring programs (Buchmueller and Carey 2018; Kaestner and Ziedan 2019; Meara et al. 2016; Kilby 2015) and pill mill laws (Lyapustina et al. 2016; Rutkow et al. 2015) which target the availability of prescription opioids for potential addicts. More recently, states have enacted opioid prescribing limits that specifically target opioid naive patients and are explicitly intended to limit transitions to addiction (Sacks et al. 2021).

Figure 9 investigates the different time paths of the effects of these policies by considering the yearly impact of counterfactual state policies in which place-based addiction transitions and place-based availability parameters above the 25th percentile are reduced to that threshold. We find similar trade-offs to those from the first exercise. In the first year of implementation, average abuse rates are 20% lower when availability is targeted compare to addiction. By the final year of our study period, however, this relationship is reversed; average abuse rates are 20% lower when addiction is targeted. The impact of policies targeting place-based addiction transitions is roughly twice as large over our 10-year study period, but policies that affect place-based availability for addicts have a much larger initial effect, suggesting important policy tradeoffs in terms of near-term compared to longer-term impacts. These findings also suggest that, as emphasized by Cutler and Glaeser (forthcoming), the stock of existing addicts creates challenges to ending the opioid epidemic.

6 Conclusion

Focusing on individuals enrolled in the Social Security Disability Insurance (SSDI) program, a population hit particularly hard by the opioid epidemic, we have explored the role of different factors in driving the prescription opioid epidemic in the decade between 2006 and 2015. The descriptive findings and the estimates from a dynamic model of opioid abuse highlight important roles for both person-specific factors and place-based factors, as well as the distinct role played by two separate place-based channels: an addiction channel which affects the rate at which people first become addicted to opioids, and an availability channel which affects the ease with which existing addicts can abuse opioids.

We estimate that equalizing place-based factors would have reduced the geographic variation in opioid abuse by about 50 percent in our study population, over our study period. Distinguishing between different place-based channels, our results suggest that supply-side policies that restrict opioid availability among existing users—such as the 2018 Medicare policy to no longer reimburse for high-dose long-term opioid prescriptions (Hoffman 2018; CMS 2018)—could have an immediate impact on rates of prescription opioid abuse. However, while most of the policy attention has focused on restricting the supply of opioids to the

already addicted, our findings suggest that changing the propensity of physicians to treat the pain of opioid naive patients with opioids—such as opioid prescribing limits specifically targeted to the opioid naive (Sacks et al. 2021)—appear to be quantitatively more important in reducing opioid abuse over a longer horizon. For example, we estimate that over our 10-year study period, reducing the upper distribution of states in their place-based effects on transitions to addiction would reduce opioid abuse by twice as much as reducing the upper distribution of states in terms of their opioid availability, although this comparison is reversed in the first few years.

The feasibility and desirability of policies targeting these place-based channels remains an important open question. In terms of the feasibility of targeting the addiction channel, recent work has documented substantial variation in physician propensity to prescribe opioids for what they believe to be legitimate reasons (e.g., Barnett et al. 2017; Schnell and Currie 2018; Eichmeyer and Zhang 2021), which is consistent with the broader evidence of substantial variation across individual physicians in other aspects of practice style (e.g. Currie and MacLeod 2020; Chan et al. 2020; Fadlon and Van Parys 2019; Van Parys 2016; Epstein and Nicholson 2009). A greater understanding of the roles that physician training (Schnell and Currie 2018), physician beliefs (Doctor et al. 2018), and organizational factors play in contributing to these differences remains an important area for further work, as is further exploration of the tradeoffs created by these alternative practice styles. In terms of the desirability of targeting the availability channel, the potential for this to cause substitution to illegal forms of opioids (Alpert et al. 2018; Evans et al. 2019) is an important area for further study.

References

Aliprantis, D., Fee, K., and Schweitzer, M. E. (2019). Opioids and the Labor Market. *Federal Reserve Bank of Cleveland, Working Paper No. 18-07R*.

Alpert, A., Evans, W. N., Lieber, E. M. J., and Powell, D. (2019). Origins of the opioid crisis and its enduring impacts. *NBER Working Paper No. 26500*.

Alpert, A., Powell, D., and Pacula, R. L. (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, D. H. (2015). The Unsustainable Rise of the Disability Rolls in the United States: Causes, Consequences, and Policy Options. In Scholz, J. K., Moon, H., and Lee, S.-H., editors, *Social Policies in an Age of Austerity: A Comparative Analysis of the US and Korea*, pages 107–136. Edward Elgar Publishing, Northampton MA.

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

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

Bohnert, A. S. B., Valenstein, M., Bair, M. J., Ganoczy, D., McCarthy, J. F., Ilgen, M. A., and Blow, F. C. (2011). Association between opioid prescribing patterns and opioid overdose-related deaths. *JAMA*, 305(13):1315–1321.

Braden, J. B., Russo, J., Fan, M.-Y., Edlund, M. J., Martin, B. C., DeVries, A., and Sullivan, M. D. (2010). Emergency department visits among recipients of chronic opioid therapy. *Archives of Internal Medicine*, 170(16):1425–1432.

Brat, G. A., Agniel, D., Beam, A., Yorkgitis, B., Bicket, M., Homer, M., Fox, K. P., Knecht, D. B., McMahill-Walraven, C. N., Palmer, N., and Kohane, I. (2018). Postsurgical prescriptions for opioid naive patients and association with overdose and misuse: retrospective cohort study. *BMJ*, 360.

Bronnenberg, B. J., Dubé, J.-P. H., and Gentzkow, M. (2012). The Evolution of Brand Preferences: Evidence from Consumer Migration. *American Economic Review*, 102(6):2472–2508.

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

Callaway, B. and Sant'Anna, P. H. (2020). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, page S0304407620303948.

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

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

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

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

Case, A. and Deaton, A. (2020). *Deaths of despair and the future of capitalism*. Princeton University Press, Princeton.

CDC (2017a). Opioid Overdose. `https://www.cdc.gov/drugoverdose/index.html`.

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

Cepeda, M. S., Fife, D., Chow, W., Mastrogiovanni, G., and Henderson, S. C. (2012). Assessing opioid shopping behaviour. *Drug Safety*, 35(4):325–334.

Chan, D., Gentzkow, M., and Yu, C. (2020). Selection with variation in diagnostic skill: Evidence from radiologists. *NBER Working Paper No. 26467*.

Charles, K. K., Hurst, E., and Schwartz, M. (2019). The Transformation of Manufacturing and the Decline in US Employment. *NBER Macroeconomics Annual*, 33:307–372.

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

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

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

Ciesielski, T., Iyengar, R., Bothra, A., Tomala, D., Cislo, G., and Gage, B. F. (2017). A tool to assess risk of de novo opioid abuse or dependence. *The American Journal of Medicine*, 129(7):699–705.

CMS (2018). 2019 Medicare Advantage and Part D Rate Announcement and Call Letter. `https://www.cms.gov/Newsroom/MediaReleaseDatabase/Fact-sheets/2018-Fact-sheets-items/2018-04-02-2.html`.

Compton, W. M., Jones, C. M., and Baldwin, G. T. (2016). Relationship between Nonmedical Prescription-Opioid Use and Heroin Use. *New England Journal of Medicine*, 374(2):154–163.

Cubanski, J., Neuman, T., and Damico, A. (2016). *Medicare's Role for People Under Age 65 with Disabilities*. Kaiser Family Foundation Report.

Currie, J., Jin, J., and Schnell, M. (2019). US Employment and Opioids: Is There a Connection? In *Health and Labor Markets*, pages 253–280. Emerald Publishing Limited.

Currie, J. and MacLeod, W. B. (2020). Understanding doctor decision making: The case of depression treatment. *Econometrica*, 88(3):847–878.

Currie, J. and Schwandt, H. (2020). The opioid epidemic was not caused by economic distress but by factors that could be more rapidly addressed. *NBER Working Paper No. 27544*.

Cutler, D. M. and Glaeser, E. L. (forthcoming). When innovation goes wrong: Technological regress and the opioid epidemic. *Journal of Economic Perspectives*.

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

Dean, A. and Kimmel, S. (2019). Free trade and opioid overdose death in the United States. *SSM - Population Health*, 8:100409.

Deyo, R. A., Hallvik, S. E., Hildebran, C., Marino, M., Dexter, E., Irvine, J. M., O'Kane, N., Van Otterloo, J., Wright, D. A., Leichtling, G., and Millet, L. M. (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.

Doctor, J. N., Nguyen, A., Lev, R., Lucas, J., Knight, T., Zhao, H., and Menchine, M. (2018). Opioid prescribing decreases after learning of a patient's fatal overdose. *Science*, 361(6402):588–590.

Doleac, J. L. and Mukherjee, A. (2018). The Moral Hazard of Lifesaving Innovations: Naloxone Access, Opioid Abuse, and Crime. *IZA Discussion Paper No. 11489*.

Dunn, K. M., Saunders, K. W., Rutter, C. M., Banta-Green, C. J., Merrill, J. O., Sullivan, M. D., Weisner, C. M., Silverberg, M. J., Campbell, C. I., Psaty, B. M., and Von Korff, M. (2010). Overdose and prescribed opioids: Associations among chronic non-cancer pain patients. *Annals of Internal Medicine*, 152(2):85–92.

Edlund, M. J., Martin, B. C., Russo, J. E., DeVries, A., Braden, J. B., and Sullivan, M. D. (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, S. and Zhang, J. (2021). Pathways Into Opioid Addiction: Evidence From Practice Variation in Emergency Departments. `https://www.jonathanzhang.net/opioid_ed`. Accessed on 2020-03-31.

Epstein, A. and Nicholson, S. (2009). The formation and evolution of physician treatment styles: An application to cesarean sections. *Journal of Health Economics*, 28:1126–1140.

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

Fadlon, I. and Van Parys, J. (2019). Primary care physician practice styles and patient care: Evidence from physician exits in medicare. *NBER Working Paper No. 26269*.

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

Fischer, B., Lusted, A., Roerecke, M., Taylor, B., and Rehm, J. (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.

GAO (1991). *The Crack Cocaine Epidemic: Health Consequences and Treatment*.

Guy, G. P., Zhang, K., Bohm, M. K., Losby, J., Lewis, B., Young, R., Murphy, L. B., and Dowell, D. (2017). Vital Signs: Changes in Opioid Prescribing in the United States, 2006-2015. *MMWR. Morbidity and Mortality Weekly Report*, 66.

Hall, A. J., Logan, J. E., Toblin, R. L., Kaplan, J. A., Kraner, J. C., Bixler, D., Crosby, A. E., and Paulozzi, L. J. (2008). Patterns of abuse among unintentional pharmaceutical overdose fatalities. *JAMA*, 300(22):2613–2620.

Harris, M. C., Kessler, L. M., Murray, M. N., and Glenn, B. (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.

Hoffman, J. (2018). Medicare Is Cracking Down on Opioids. Doctors Fear Pain Patients Will Suffer. *The New York Times*.

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

Ives, T. J., Chelminski, P. R., Hammett-Stabler, C. A., Malone, R. M., Perhac, J. S., Potisek, N. M., Shilliday, B. B., DeWalt, D. A., and Pignone, M. P. (2006). Predictors of opioid misuse in patients with chronic pain: a prospective cohort study. *BMC Health Services Research*, 6:46.

Jeffery, M. M., Hooten, W. M., Hess, E. P., Meara, E. R., Ross, J. S., Henk, H. J., Borgundvaag, B., Shah, N. D., and Bellolio, M. F. (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, A. B., Goldman, D., Weaver, L., and Karaca-Mandic, P. (2014). Opioid prescribing by multiple providers in Medicare: retrospective observational study of insurance claims. *BMJ*, 348:g1393.

Kaestner, R. and Ziedan, E. (2019). Mortality and socioeconomic consequences of prescription opioids: Evidence from state policies. *NBER Working Paper No. 26135*.

Katz, J. (2017). Short Answers to Hard Questions About the Opioid Crisis. *The New York Times*.

Khan, N. F., Bateman, B. T., Landon, J. E., and Gagne, J. J. (2019). Association of opioid overdose with opioid prescriptions to family members. *JAMA Internal Medicine*, 179(9):1186–1192.

Kilby, A. 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, J., Gorfinkel, L., Fairbairn, N., Amato, L., Ahamad, K., Nolan, S., Simel, D. L., and Wood, E. (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, K. D., Murphy, S. L., Xu, J., and Arias, E. (2019). Deaths: Final data for 2017. *National Vital Statistics Reports*, 68(9):1–76.

Krueger, A. 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.

Laird, J. and Nielsen, T. (2016). The Effects of Physician Prescribing Behaviors on Prescription Drug Use and Labor Supply: Evidence from Movers in Denmark. `https://scholar.harvard.edu/files/lairdja/files/Laird_JMP.pdf`. Accessed on 2021-03-31.

Larrick, A. (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, J., Liu, Y., Paulozzi, L., Zhang, K., and Jones, C. (2013). Opioid Prescribing in Emergency Departments: The Prevalence of Potentially Inappropriate Prescribing and Misuse. *Medical Care*, 51(8):646–653.

Lyapustina, T., Rutkow, L., Chang, H.-Y., Daubresse, M., Ramji, A. F., Faul, M., Stuart, E. A., and Alexander, G. C. (2016). Effect of a pill mill law on opioid prescribing and utilization: The case of Texas. *Drug and Alcohol Dependence*, 159:190–197.

Maclean, J. C., Mallatt, J., Ruhm, C. J., and Simon, K. (2020). Economic studies on the opioid crisis: A review. *NBER Working Paper No. 28067*.

McDonald, D. C., Carlson, K., and Izrael, D. (2012). Geographic variation in opioid prescribing in the U.S. *The Journal of Pain*, 13(10):988–996.

Meara, E., Horwitz, J. R., Powell, W., McClelland, L., Zhou, W., O'Malley, A. J., and Morden, N. E. (2016). State Legal Restrictions and Prescription-Opioid Use among Disabled Adults. *New England Journal of Medicine*, 375(1):44–53.

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

Morden, N. E., Munson, J. C., Colla, C. H., Skinner, J. S., Bynum, J. P. W., Zhou, W., and Meara, E. (2014). Prescription opioid use among disabled Medicare beneficiaries: intensity, trends, and regional variation. *Medical Care*, 52(9):852–859.

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

Muhuri, P., Gfroerer, J., and Davies, C. (2013). Associations of Nonmedical Pain Reliever Use and Initiation of Heroin Use in the United States. *CBHSQ Data Review*.

NCHS (2021). *CDC WONDER: National Drug Overdose (OD) Deaths, 1999-2019*.

Parente, S. T., Kim, S. S., Finch, M. D., Schloff, L. A., Rector, T. S., Seifeldin, R., and Haddox, J. D. (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, L. J., Kilbourne, E. M., Shah, N. G., Nolte, K. B., Desai, H. A., Landen, M. G., Harvey, W., and Loring, L. D. (2012). A History of Being Prescribed Controlled Substances and Risk of Drug Overdose Death. *Pain Medicine*, 13(1):87–95.

Powell, D., Pacula, R. L., and Taylor, E. (2020). How increasing medical access to opioids contributes to the opioid epidemic: Evidence from Medicare Part D. *Journal of Health Economics*, 71:102286.

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

Rees, D. I., Sabia, J. J., Argys, L. M., Dave, D., and Latshaw, J. (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. B., White, A. G., Birnbaum, H. G., Schiller, M., Brown, D. A., and Roland, C. L. (2012). A model to identify patients at risk for prescription opioid abuse, dependence, and misuse. *Pain Medicine*, 13(9):1162–1173.

Rough, K., Huybrechts, K. F., Hernandez-Diaz, S., Desai, R. J., Patorno, E., and Bateman, B. T. (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].

Ruhm, C. J. (2018). Deaths of Despair or Drug Problems? *NBER Working Paper No. 24188*.

Rummans, T. A., Burton, M. C., and Dawson, N. L. (2018). How Good Intentions Contributed to Bad Outcomes: The Opioid Crisis. *Mayo Clinic Proceedings*, 93(3):344–350.

Rutkow, L., Chang, H.-Y., Daubresse, M., Webster, D. W., Stuart, E. A., and Alexander, G. C. (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, D. W., Hollingsworth, A., Nguyen, T., and Simon, K. (2021). Can policy affect initiation of addictive substance use? Evidence from opioid prescribing. *Journal of Health Economics*, 76:102397.

Schieber, L. Z., Guy, Gery P., J., Seth, P., Young, R., Mattson, C. L., Mikosz, C. A., and Schieber, R. A. (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, M. and Currie, J. (2018). Addressing the opioid epidemic: Is there a role for physician education? *American Journal of Health Economics*, pages 383–410.

Shah, A., Hayes, C. J., and Martin, B. C. (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, J. (2011). Chapter Two - Causes and Consequences of Regional Variations in Health. In Pauly, M. V., Mcguire, T. G., and Barros, P. P., editors, *Handbook of Health Economics*, volume 2 of *Handbook of Health Economics*, pages 45–93. Elsevier.

Song, Y., Skinner, J., Bynum, J., Sutherland, J., Wennberg, J. E., and Fisher, E. S. (2010). Regional variations in diagnostic practices. *New England Journal of Medicine*, 363(1):45–53.

SSA (2017). *Annual Statistical Supplement to the Social Security Bulletin, 2016*.

Sullivan, M. D., Edlund, M. J., Fan, M.-Y., Devries, A., Brennan Braden, J., and Martin, B. C. (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, E. C., Darnall, B. D., Baker, L. C., and Mackey, S. (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, L. and Abraham, S. (2020). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, page S030440762030378X.

Turk, D. C., Swanson, K. S., and Gatchel, R. J. (2008). Predicting opioid misuse by chronic pain patients: a systematic review and literature synthesis. *The Clinical Journal of Pain*, 24(6):497–508.

Van Parys, J. (2016). Variation in physician practice styles within and across emergency departments. *PloS One*, 11(8).

Venkataramani, A. S., Bair, E. F., O'Brien, R. L., and Tsai, A. C. (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, L. R. (2017). Risk Factors for Opioid-Use Disorder and Overdose. *Anesthesia & Analgesia*, 125(5):1741–1748.

Wei, Y.-J. J., Chen, C., Sarayani, A., and Winterstein, A. G. (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, A. G., Birnbaum, H. G., Schiller, M., Tang, J., and Katz, N. P. (2009). Analytic models to identify patients at risk for prescription opioid abuse. *The American Journal of Managed Care*, 15(12):897–906.

Yang, Z., Wilsey, B., Bohm, M., Weyrich, M., Roy, K., Ritley, D., Jones, C., and Melnikow, J. (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 Opioid Abuse

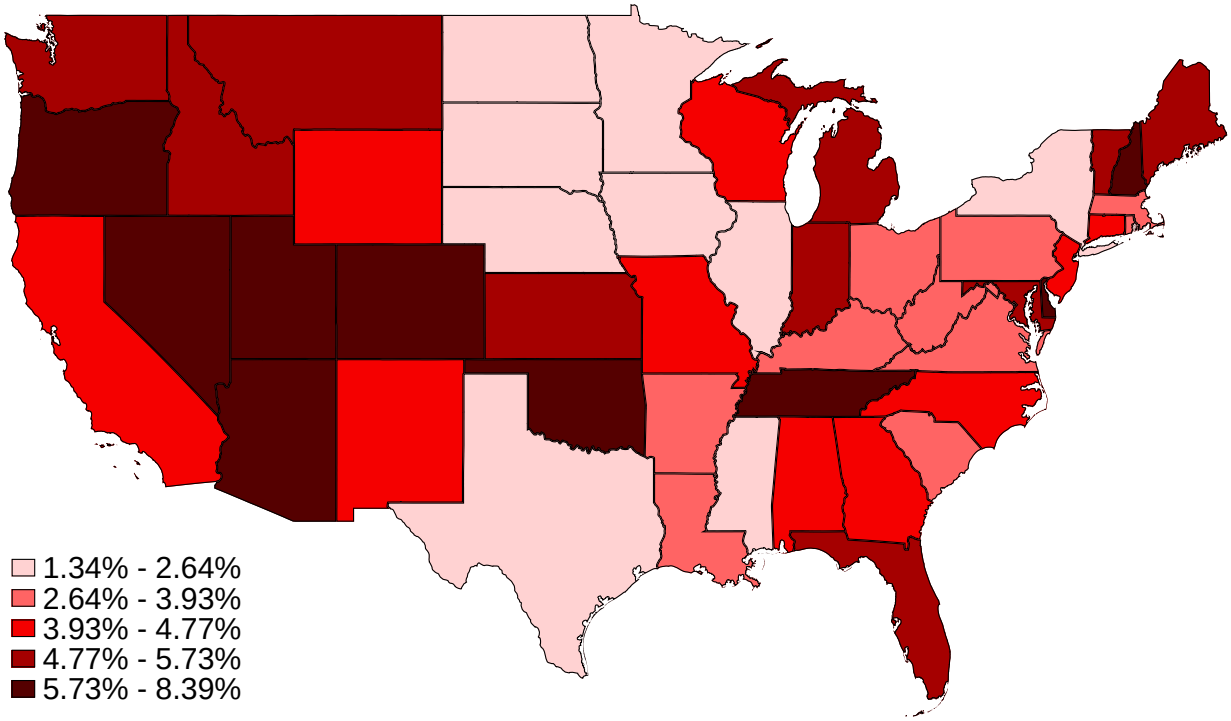

Notes: Figure reports state-level averages for the rate of opioid abuse among all non-mover years (N = 2,325,094 enrollee-years).

Figure 2: Interpreting Event Study Specification

(a) Case 1: Only Person Effects Vary

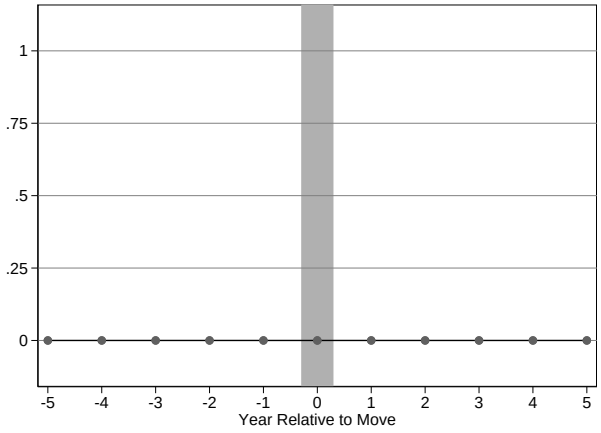

(b) Case 2: Only Availability Effects Vary

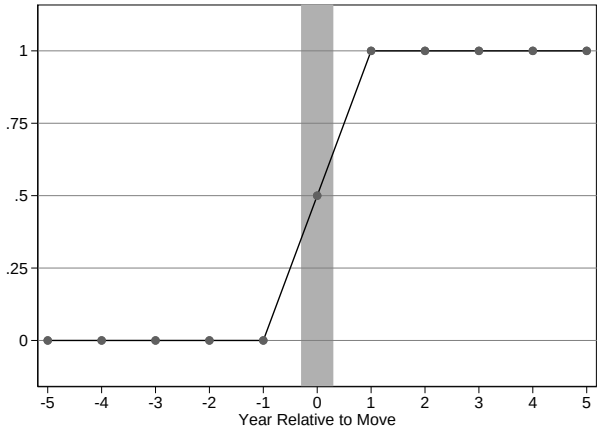

(c) Case 3: Both Person and Availability Effects Vary

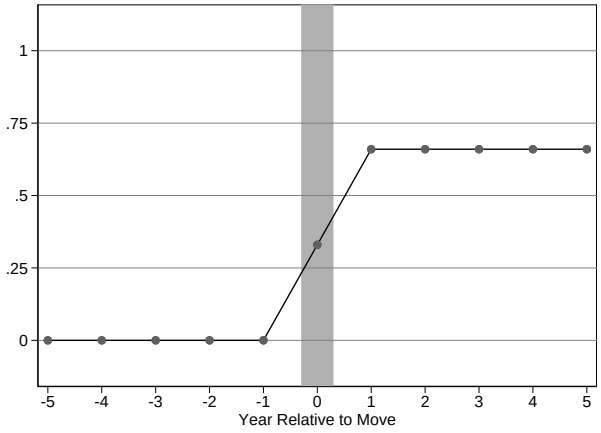

(d) Case 4: All Channels Vary

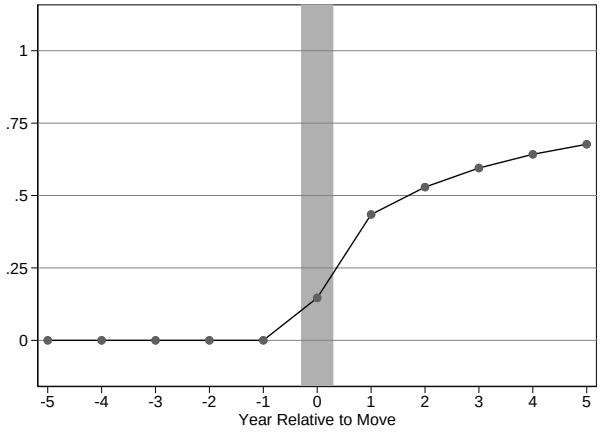

Notes: Figure shows the patterns of event study coefficients $\mu_{r(i,t)}$—the coefficient on year relative to move interacted with $\hat{\delta}_{ct}$—as calculated according to equation (5) in a theoretical setting with a single cohort moving between two locations and the coefficients plotted on the y-axis. In this setting, abuse rates are determined according to the parameters described in Section 3, and we chose an illustrative example set of parameters for each example. 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 the 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. Finally, in panel D, the two locations differ in their person-based factors, availability effects, and their place-based effects on transitions to addiction.

Figure 3: Distribution of Origin - Destination Differences in Opioid Abuse Rates

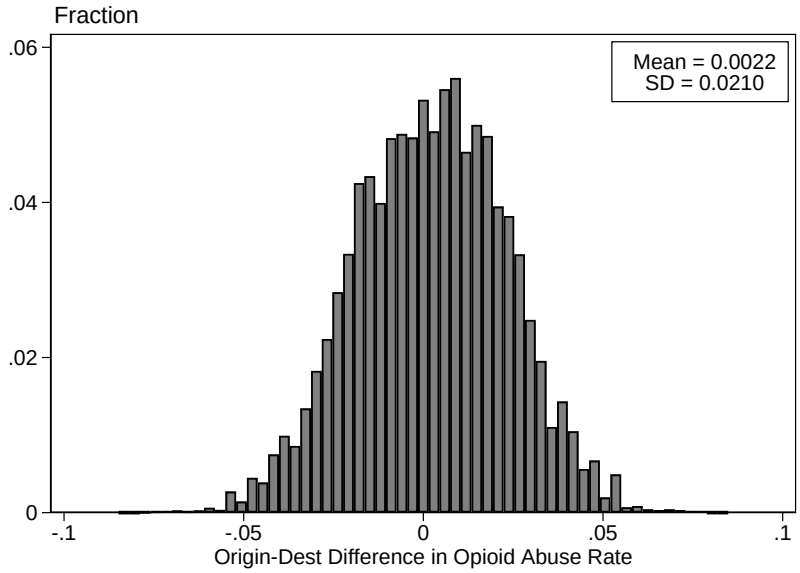

Notes: Figure shows the distribution across movers of the difference in the average opioid abuse rates between their origin and destination states in a given year ($\hat{\delta}_{ct}$) The sample is all movers (N = 90,890 enrollees).

Figure 4: Event Study

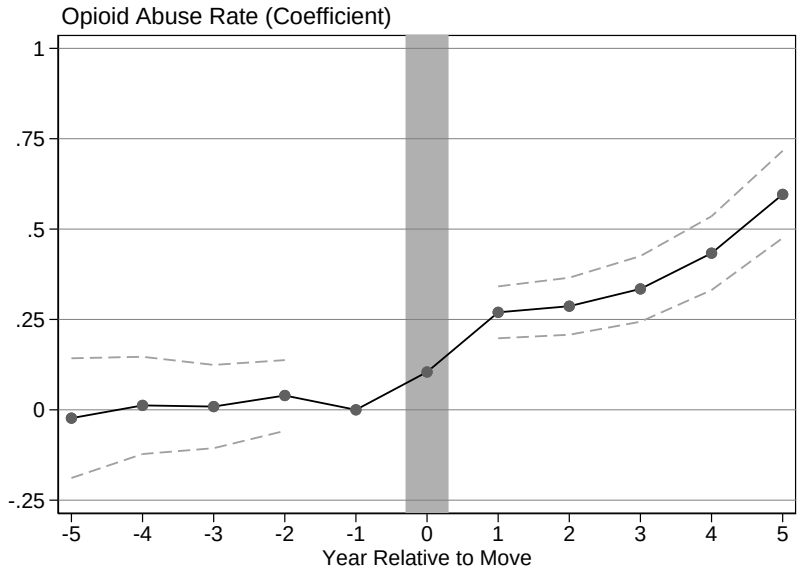

Notes: Figure shows the weighted average across move years of estimates from equation (3) of $\mu_{r(i,t)}$—the coefficient on year relative to move interacted with $\hat{\delta}_{ct}$—where the coefficient for relative year -1 is normalized to 0 and weights are given by sample size. The dependent variable $\tilde{y}_{ict}$ is the residual of the mover's abuse rate after subtracting the predicted abuse rate using origin-year and five-year age bin fixed effects as estimated from non-movers. The dashed lines are upper and lower bounds of the 95% confidence interval. We construct this confidence interval following a three-step procedure. In the first step, for each state-year, we construct the asymptotic distribution of its average rate of abuse $\bar{y}_{jt}$ according to the central limit theorem, which is a normal distribution with mean and standard deviation calculated from the data. In the second step, we bootstrap equation (1) with 50 repetitions drawn at the enrollee-year level, making a random draw from the asymptotic distribution for each mover's origin, destination, and year to construct their $\hat{\delta}_{it}$ for each repetition. Finally, using these bootstraps to estimate the standard errors for each move year estimate, we calculate the variance of the weighted average coefficients by assuming that the move year coefficients are independent. The sample is all mover-years (N = 521,523 enrollee-years).

Figure 5: Change in Abuse Rate by Size of Move

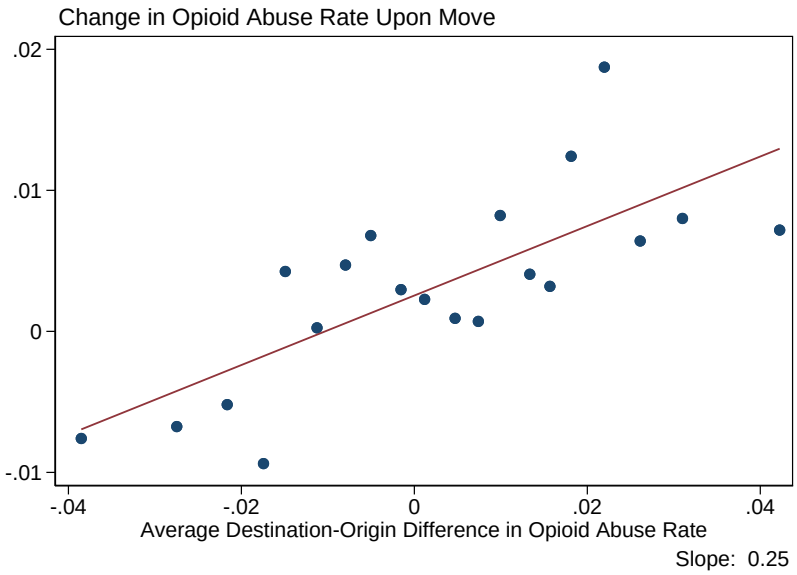

Notes: Figure shows changes in the rates of opioid abuse before and after move. For each mover, we calculate the average difference $\hat{\delta}_{ct}$ in the rate of opioid abuse between their origin and destination states for all years they are observed between 1-5 years pre-move and 1-5 years post-move. We then group these average differences—$\bar{\hat{\delta}}_{ct}$—into ventiles. The x-axis displays the mean of $\bar{\hat{\delta}}_{it}$ for movers in each ventile. The y-axis shows, for each ventile, the average opioid abuse rate 1-5 years post-move minus the average opioid abuse rate 1-5 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. The sample is all mover-years between one and five years pre-move and one and five years post-move (N = 451,124 enrollee-years).

Figure 6: Event Studies - Naives and Prior Users

(a) Naives

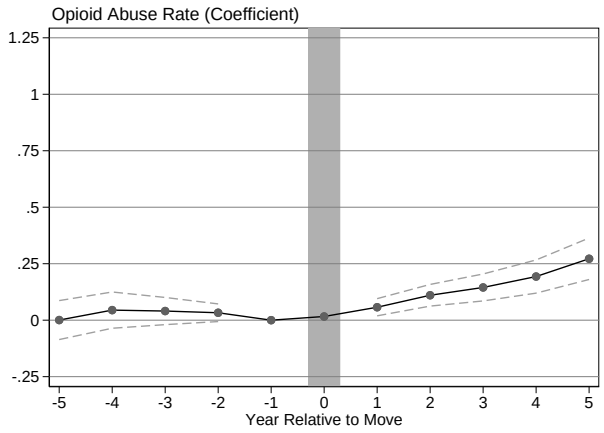

(b) Prior Users

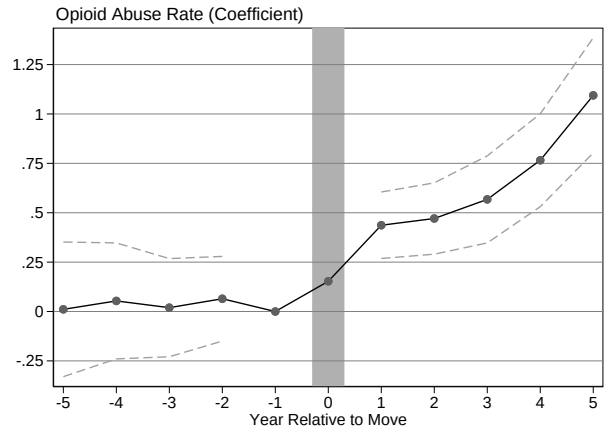

Notes: Figure shows the weighted average across move years of estimates from equation (3) of $\mu_{r(i,t)}$—the coefficient on year relative to move interacted with $\hat{\delta}_{ct}$—where the coefficient for relative year -1 is normalized to 0 and weights are given by sample size. We show these estimates separately for opioid naive ("naive") and prior users, where the coefficient for relative year -1 is normalized to 0. "Naives" are those with no opioid use in relative year -1, while prior users filled at least one opioid prescription in that year. We omit the approximately 20% of enrollee-years for enrollees who are not observed in the year before moving. The dependent variable $\tilde{y}_{ict}$ is the residual of the mover's abuse rate after subtracting the predicted abuse rate using origin-year and five-year age bin fixed effects as estimated with non-movers with the corresponding opioid use in the same calendar year as represented by relative year -1. The dashed lines are upper and lower bounds of the 95% confidence interval. We construct this confidence interval following a three-step procedure. In the first step, for each state-year, we construct the asymptotic distribution of its average rate of abuse $\bar{y}_{jt}$ according to the central limit theorem, which is a normal distribution with mean and standard deviation calculated from the data. In the second step, we bootstrap equation (3) with 50 repetitions drawn at the enrollee-year level, making a random draw from the asymptotic distribution for each mover's origin, destination, and year to construct their $\hat{\delta}_{ct}$ for each repetition. Finally, using these bootstraps to estimate the standard errors for each move year estimate, we calculate the variance of the weighted average coefficients assuming that the move year coefficients are independent. The samples are 219,100 enrollee-years (naives) and 202,309 enrollee-years (prior users).

Figure 7: Event Study in Simulation vs. Data

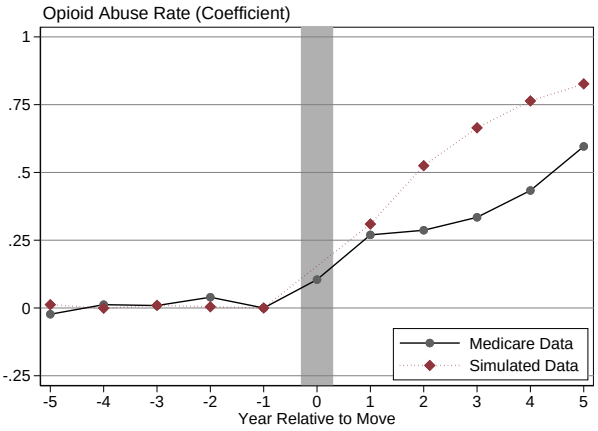

Notes: Figure shows the weighted average across move years of estimates from equation (3) of $\mu_{r(i,t)}$—the coefficient on year relative to move interacted with $\hat{\delta}_{ct}$—as estimated on the simulated data and the Medicare data, where the coefficient for relative year -1 is normalized to 0 and weights are given by sample size. The red dashed line shows the aggregate coefficients estimated from our simulated event study, where the yearly rates of opioid abuse among movers and non-movers are simulated according to the data-generating process described in Section 5.2. The solid black line is identical to Figure (4) and plots our aggregate event study coefficients as estimated on the Medicare data.

Figure 8: Changes in Geographic Variation in Abuse Relative to Baseline

(a) Equalizing Person-Based Factors

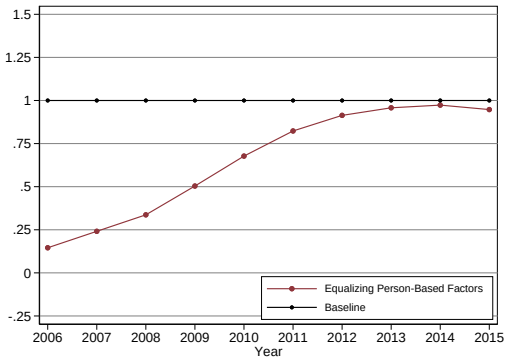

(b) Equalizing Place-Based Factors

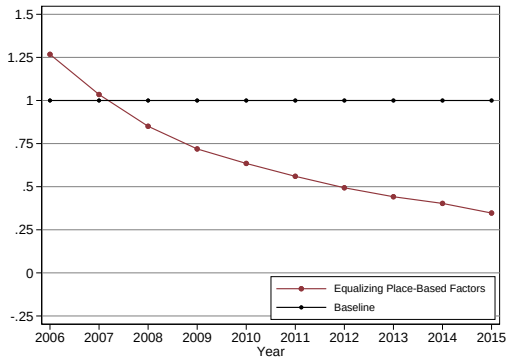

(c) Equalizing Initial Addiction States

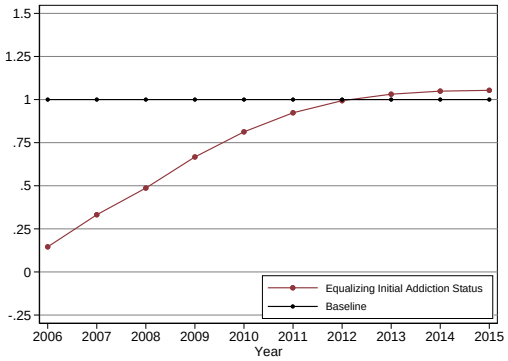

(d) Equalizing Addiction Transition Probabilities

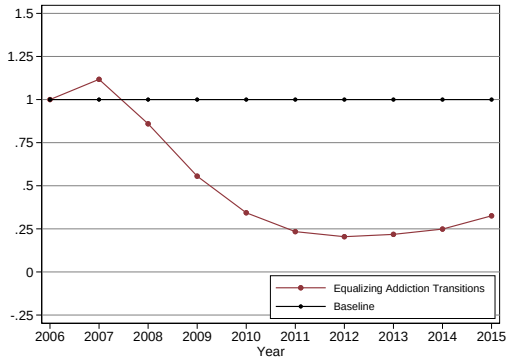

(e) Equalizing Availability Channel

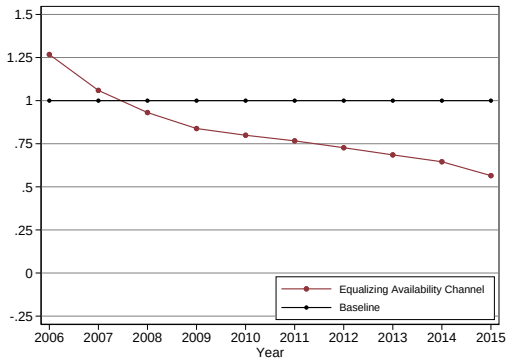

Notes: In this figure, we set all the relevant parameters equal to their median values and plot the yearly variation in abuse relative to the baseline simulation on the y-axis. This is defined as the ratio of the variance of abuse rates across states in the counterfactual to the variance of abuse rates across states in the baseline simulation. All objects are computed using average state-level abuse rates and assign equal weight to each state. Place-based parameters refer to place-based addiction π_j as well as availability γ_{jt}. Person-based parameters refer to person-based addiction η_i as well as initial addiction state a_{i0}. Initial addiction state refers to exclusively a_{i0}, while addiction transition parameters refer to π_j and η_i. Finally, availability parameters refer to γ_{jt}. For the parameter a_{i0}, we directly compute the median values among states. For γ_{jt}, we set this parameter equal to the median value among states in each year. For our addiction transition parameters, the median values refer to the median steady state and convergence speed implied by the combination of parameters. As introduced in Section 5, the steady state is defined as $\frac{\pi_j^+ + \eta_j^+}{\pi_j^+ + \pi_j^- + \eta_j^+ + \eta_j^-}$ and the convergence speed is defined as $\pi_j^+ + \pi_j^- + \eta_j^+ + \eta_j^-$ where $\eta_i = \eta_j \; \forall \; i \in \mathscr{I}_j$. For each subset of addiction transition parameters—place-based addiction transitions and person-based addiction transitions—we define the steady state and convergence speed holding the other set of parameters constant.

Figure 9: Comparison of Counterfactual Place-Based Policies

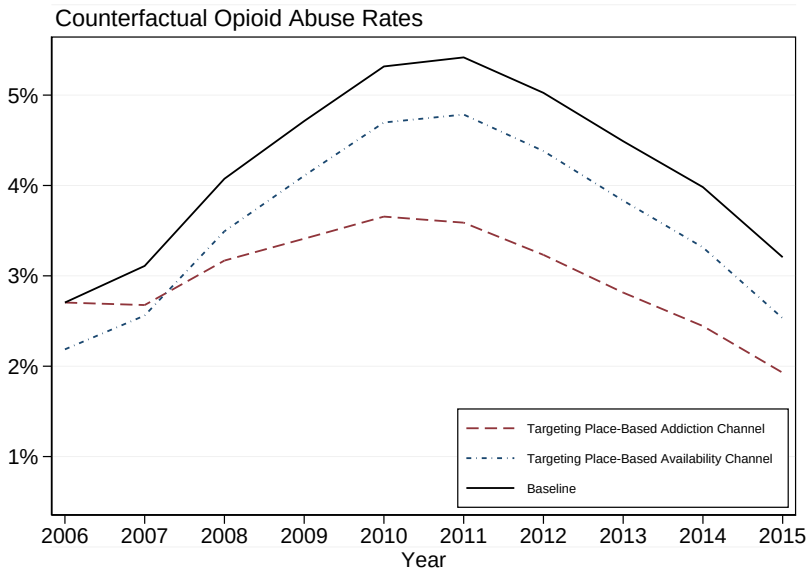

Notes: In this figure, we show the counterfactual average yearly abuse rates under the baseline simulation and two counterfactual policies. The average yearly abuse rate is computed as a weighted mean according to the number of non-mover observations in the Medicare data. In the place-based availability policy, we reduce place-based availability parameters (γ_{jt}) in each year above the 25th percentile to the 25th percentile value in that year. Likewise in the place-based addiction policy, we reduce the place-based addiction parameters above the 25th percentile to the 25th percentile value. For place-based addiction parameters, we define percentiles by the steady state and convergence speed, as defined in Section 5, that are implied by the combination of place-based addiction transition parameters while holding person-based addiction parameters constant. Thus, we choose the percentile values of place-based addiction parameters consistent with the percentile values of $\frac{\pi_j^+}{\pi_j^+ + \pi_j^-}$ and $\pi_j^+ + \pi_j^-$.

Table 1: Summary Statistics

	(1) Movers	(2) Non-movers
Female	56%	52%
White	74%	72%
Medicaid	58%	57%
Age:		
< 40	13%	11%
40 - 60	46%	44%
> 60	41%	45%
Average age	56.2	57.7
Region:		
Northeast	15%	19%
West	22%	18%
Midwest	18%	22%
South	44%	41%
Opioid Use:		
Any Opioids	47.1%	41.8%
Prescriptions in year before move ("prior user")	38.8%	
No prescriptions in year before move ("opioid naive")	42.0%	
No observation in year before move	19.2%	
Abuse	4.8%	4.1%
Number of enrollee-years	521,523	13,349,773
Number of enrollees	90,890	2,325,094

Notes: All rows except for enrollee-years and enrollees report the share of enrollees or enrollee-years within the given population with the indicated characteristic. "Any Opioids" and "Opioid Abuse" are averaged over all enrollee-years, while all other statistics are averaged at the enrollee-level, with "Region," "Medicare," and "Age" defined in the year before move ($r = -1$).

Table 2: Event Study Coefficients - Opioid Abuse Rates

	(1)	(2)	(3)
	All	Naive	Prior User
1 year post-move	0.270	0.057	0.437
	(0.037)	(0.019)	(0.086)
5 years post-move	0.596	0.272	1.094
	(0.029)	(0.047)	(0.149)
Enrollees	90,890	32,648	30,460
Enrollee-years	521,523	219,100	202,309
Average abuse rate 1 year pre-move	0.048	0.000	0.100

Notes: Table reports the coefficients and their bootstrapped standard errors (in parentheses) in relative year 1 and relative year 5 for the baseline sample ("All"), naive enrollees, and prior users with opioid abuse as the outcome variable. Estimates correspond to the weighted average across move years of estimates from equation (3) of $\mu_{r(i,t)}$—the coefficient on year relative to move interacted with $\hat{\delta}_{ct}$—where the coefficient for relative year -1 is normalized to 0 and weights are given by sample size. 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. Except in the baseline sample, we omit the approximately 20% of enrollee-years for enrollees who are not observed in the year before moving. Standard errors are calculated following a three-step procedure. In the first step, for each state-year, we construct the asymptotic distribution of its average rate of abuse $\bar{y}_{jt}$ according to the central limit theorem, which is a normal distribution with mean and standard deviation calculated from the data. In the second step, we bootstrap equation (3) with 50 repetitions drawn at the enrollee-year level, making a random draw from the asymptotic distribution for each mover's origin, destination, and year to construct the $\hat{\delta}_{ct}$ for each repetition. Finally, using these bootstraps to estimate the standard errors for each move year estimate, we calculate the standard deviation of the weighted average coefficients by assuming that the move year coefficients are independent.

Table 3: Abuse Decompositions

	Share variance reduced when we equalize	
(1)	All Place-Based Factors (π_j, γ_{jt})	47%
(2)	All Person-Based Factors (η_i, a_{i0})	28%
(3)	Initial Addiction States (a_{i0})	14%
(4)	Addiction Transition Probabilities (π_j, η_i)	63%
(5)	Availability (γ_{jt})	26%

Notes: All objects are computed using average state-level abuse rates and assign equal weight to each state. For all rows, we compute the variance of simulated opioid abuse rates while equalizing the relevant structural parameters by setting them equal to their "median value" and present the effect of these counterfactuals on the simulated variance. Place-based parameters refer to place-based addiction π_j as well as availability γ_{jt}. Person-based parameters refer to person-based addiction η_i as well as initial addiction state a_{i0}. Initial addiction state refers to a_{i0} while addiction transition probabilities refer to π_j and η_i. Finally, availability parameters refer to γ_{jt}. For the parameter a_{i0}, we directly compute the median values among states. For γ_{jt}, we set this parameter equal to the median value among states in each year. For our place-based addiction parameters, we choose the median steady state and convergence speed. For our addiction transition parameters, the median values refer to the median steady state and convergence speed implied by the combination of parameters. As introduced in Section 5, under the simplifying assumption that all individuals from a location share the same addiction propensities, the steady state for non-movers is defined as $\frac{\pi_j^+ + \eta_j^+}{\pi_j^+ + \pi_j^- + \eta_j^+ + \eta_j^-}$ and the convergence speed is defined as $\pi_j^+ + \pi_j^- + \eta_j^+ + \eta_j^-$ where $\eta_i = \eta_j \ \forall \ i \in \mathscr{I}_j$. For each subset of addiction transition parameters—place-based addiction transitions and person-based addiction transitions—we define the steady state and convergence speed holding the other subset of addiction transition parameters constant.

Appendix A: Defining opioid abuse

Measuring opioid abuse 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 opioid abuse from claims data. However, the literature uses several proxies for likely opioid abuse 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 more detailed 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 MED (usually above 100-120 mg) are at higher risk for diagnoses of opioid dependence (Sullivan et al. 2010; Edlund et al. 2014), ER visits associated with opioids (Braden et al. 2010), 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 (Yang et al. 2015; Hall et al. 2008).

We construct several measures of prescription opioid use and potential abuse. 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 ingredient described in the MED conversion table (Appendix Table A.4).

We define and separately analyze three indicator variables as proxies of abuse: High MED, Many Prescribers, and Overlapping Prescriptions. Each of these proxies are defined following previous medical literature; High MED and Many Prescribers are defined following Meara et al. (2016) while Overlapping Prescriptions is defined following Jena et al. (2014) and Logan et al. (2013).

Our primary measure of abuse is "High MED". This is an indicator for whether any quarterly MED is greater than 120 mg. To define it, we compute the MED for each 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 each quarter. "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 "ran out." 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 from 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, and Chronic Opioid Use. Following Morden et al. (2014), we define "chronic opioid use" 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. Similarly, we define opioid use disorder as the set of

so-labelled diagnoses in the claims data.

Comparison across measures

We begin by noting that about four percent of our enrollee-years have a high average daily MED (High MED), six percent of enrollee-years include prescriptions from four or more unique prescribers (Many Prescribers), and about 15 percent of enrollee-years have overlapping opioid prescriptions (Overlapping Prescriptions). The pairwise rank correlations among these three proxies of opioid abuse are high (Appendix Table A.5) and all are positively correlated with our other indicators of prescription opioid use.

To compare our various proxies of opioid abuse, we limit our analysis to the non-mover sample with traditional Medicare for which we can fully observe observe opioid poisonings and opioid use disorders. We further restrict to non-mover enrollee years for which we observe the following year—about 60 percent of our enrollee-years—in order to examine the relationship between our measures of opioid use in year $t-1$ and adverse outcomes in year t (Table A.6).

In panel A, we see that on average, one percent of our sample is diagnosed with an opioid use disorder each year and two-tenths of a percent of our sample is diagnosed with an opioid poisoning event. These adverse outcomes are significantly more prevalent among individuals with our measures of opioid abuse in the year preceding. For example, while only four percent of our overall population is classified as High MED, this measure is about seven times more common among enrollees with an opioid poisoning event in the following year and about eight times more common among enrollees with an opioid poisoning event in the subsequent year. We then translate these rates into a measure of diagnostic accuracy by constructing the positive likelihood ratio in panel B, defined as the ratio of true positives to false positives, for each of our proxies of opioid abuse. We choose the highest performing abuse measure—High MED—across both poisonings and opioid use disorders as our primary outcome.

Finally, for our preferred measure of opioid abuse—High MED—we examine how geographic patterns of abuse in our population correlate with geographic patterns in adverse opioid outcomes in the overall US population. To do so, we use 2006-2014 state-level data from the CDC Multiple Cause of Death (MCD) File[12] 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 2006-2014 National Survey on Drug Use and Health (NSDUH)[13] as the share of the adult population who responded "yes" to "non-medical use of pain relievers in the past year." At the state level, both measures have a high rank correlation with our opioid abuse measure; the average yearly rank correlation is 0.52 for the opioid-related death rate, and 0.46 for the self-reported opioid misuse rate.

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

[13] Available at http://datafiles.samhsa.gov/study-series/national-survey-drug-use-and-health-nsduh-nid13517.

Appendix B: 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 U.S (CDC 2017b). The aggregated QuintilesIMS data set contains the number of opioid prescriptions per capita in each year.[14] We define an identical variable in our data, the number of opioid prescription fills per capita (Opioid Fills). Both are defined over the 2006-2014 time period.

Appendix Figure A.6 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.

Substitution to illegal opioids and adverse events

We cannot directly examine illegal opioid use in our data. However, we conduct two exercises and find results that are consistent with moving to a higher prescription opioid abuse state impacting total opioid use. This of course does not preclude the existence of quantitatively important substitution from prescription opioids to illegal forms of opioids, but it is suggestive that any such substitution is not fully offsetting.

First, the existing literature has documented substitution from prescription opioids to heroin, but little evidence of substitution from heroin to prescription opioids (Muhuri et al. 2013; Compton et al. 2016). Therefore, if substitution was an important factor, we would expect to see bigger effects on prescription opioid abuse from moves down—from higher opioid abuse states to lower abuse states—as users substitute towards heroin. Conversely, we expect more muted effects on prescription opioid abuse from moves up from lower opioid abuse states to higher opioid abuse states. However, the results in Appendix Table A.8, particularly five years after the move, suggest that in the long run, moves up have a larger impact on prescription opioid abuse than moves down.

Second, we can look at whether moves across states with different rates of prescription opioid abuse are associated with changes in hospital admissions and ER visits that are related to opioids.[15] To do so, we re-estimate the event study equation (3) with the outcome defined as a "poisoning event"—specifically a non-fatal inpatient admission or ED visit related to a poisoning (based on ICD-9 diagnosis codes)—and the measure of the "size of the move"—$\hat{\delta}_{ct}$—based on differences in opioid abuse rates among non-movers. The analysis thus asks whether moving to a state with a higher rate of prescription opioid abuse is associated with a higher likelihood of a poisoning event. We use the more general "poisonings" as opposed to "drug poisonings" or "opioid poisonings" because different hospitals may differ in the specificity of their codings. On average, the data indicate that about 80 percent of poisonings are "drug poisonings" and about one-quarter of drug poisonings are are "opioid poisonings". The results—shown in Appendix Table A.9 and Appendix Figure A.7—are fairly noisy and inconclusive; presumably due to the low frequency of poisoning events.

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

[15] Opioid-related deaths are not possible to study with our approach: they are exceedingly rare, cannot be analyzed in a panel, and CMS stopped providing cause of death data after 2008.

Robustness of Descriptive Analysis

The descriptive results in Section 4 on opioid abuse are robust to several alternative specifications. Appendix Table A.7 summarizes these findings; the underlying event studies are shown in Appendix Figure A.4.

Panel A presents our baseline results for ease of comparison. These results are similar if we 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; there are about 700 commuting zones and 3,000 counties in the United States.

We also explore robustness to modifications of our event study equation. Panel D presents our results if we include individual fixed effects in equation (3), while panel E presents our results if we include relative year fixed effects. Panel F includes both individual fixed effects and relative year fixed effects in the spirit of the estimating equation from Finkelstein et al. (2016). While the inclusion of individual fixed effects seems to add slight pre-trends, the basic pattern of our event studies—a discrete jump upon move and then post-trends—holds across each of these alternative specifications. In addition, Table A.2 suggests that these features are robust to balancing our sample across event time after the move, indicating that these features are not simply driven by the changing composition of groups at each length of exposure to the new location.

The results are also stable across alternative ways of defining our sample to handle entry or exit by enrollees during our sample period. Our baseline sample restricts to enrollee-years that have 12 months of Part D coverage. A given enrollee may therefore enter or exit the sample if she gains or loses Part D coverage. Panel G therefore shows results if we further restrict the sample to enrollees who had Part D coverage in all years they appear in our sample, while panel H restricts to enrollees who are alive in all years during our study period. To further explore potential impacts of selection into or sample, Appendix Figure A.5 estimates equation (3) using as the dependent variable whether the enrollee-year had 12 months of Part D coverage and whether the enrollee died in that year. For all of these analyses $\hat{\delta}_{ct}$ is still defined based on opioid abuse. The results show little evidence of selection based on survival and some evidence of selection into Part D upon move; for example, a move from a 20th percentile opioid abuse state to a 80th percentile state is associated with about a 2 percentage point increase in the probability of having Part D coverage (relative to a mean of 80 percent). Reassuringly, however, Table A.7 shows that restricting attention to individuals whose selection along these margins does not change during the sample does not affect our main results.

Finally, we consider the robustness of our analysis to alternative ways the literature has measured opioid abuse. As we discussed in Appendix A, two of these measures of abuse include the Many Prescribers and Overlapping Prescriptions measures. We show here that the basic event study results from estimating equation 3 are similar if we instead use these alternative measures instead of our baseline measure, High MED. Appendix Table A.3 presents summary statistics on these measures, Appendix Table A.1 summarizes the results, and Appendix Figure A.3 shows the corresponding event studies. The immediate discrete jump upon move using Overlapping Prescriptions and the abuse index are slightly smaller than our baseline index result, while the jumps for Many Prescribers are slightly larger.

We note that the Many Prescribers measure does not feature the gradual post-move convergence of our other specifications. A lack of post-trends is fully plausible within our addiction framework. In Section 4.1, we analyzed the case (Case 4) in which place-based addiction transitions were higher in higher abuse locations.[16] However, if we relax this assumption—consider, for example, if high abuse rates were primarily driven by person-specific factors or availability rather than place-based addiction factors—the average abuse rates in a cohort could converge in the opposite direction of δ_{ct}. Thus, we could see zero or even negative post-trends.

[16]We define "higher" in the sense that π_j^+ is higher and π_j^- is lower so that holding person-factors constant, an individual is more likely to both become and stay addicted in the location with "higher" place-based addiction transitions.

Appendix C: Derivation of Equation (4)

To derive equation (4) from the event study regression in equation (3), observe that the coefficient for each relative year may be simply be written as the result of a regression through the origin, such that

$$\hat{\mu}_r = \frac{\sum_i \delta_{ict} \left(y_{ict} - \hat{y}_{ict}^{predict} \right)}{\sum_i \delta_{ict}^2} \quad .$$

The assumption that movers are randomly sampled from their origin population implies that $y_{ict}^{predict}$ is simply the sample analog of $E_{i \in \mathscr{I}_c} y_{it} \left(\mathbf{h}_{ct}^0 \right)$. We also note that y_{ict} is equivalent to our previously defined term $y_{ict}(\mathbf{h}_{ct})$. Decomposing our summation by cohort and replacing these terms yields

$$\begin{aligned} \hat{\mu}_r &= \frac{\sum_{c \in \mathscr{C}_m} \delta_{ct} \sum_{i \in \mathscr{I}_c} \left(y_{ict} - \hat{y}_{ict}^{predict} \right)}{\sum_i \delta_{ict}^2} \\ &= \frac{\sum_{c \in \mathscr{C}_m} \delta_{ct} \sum_{i \in \mathscr{I}_c} \left[E_{i \in \mathscr{I}_c} \left[y_{it}(\mathbf{h}_{ct}) \right] - E_{i \in \mathscr{I}_c} \left[\hat{y}_{it}(\mathbf{h}_{ct}^0) \right] \right]}{\sum_i \delta_{ict}^2}. \end{aligned}$$

Per equation (1), we denote the sample period-t average treatment effect on movers in cohort c of moving relative to remaining in their origin as $\hat{T}_{ct} \equiv E_{i \in \mathscr{I}_c} \left[y_{it}(\mathbf{h}_{ct}) - \hat{y}_{it}(\mathbf{h}_{ct}^0) \right]$. This allows us to simplify the coefficient for each relative year as

$$\hat{\mu}_r = \frac{\sum_{c \in \mathscr{C}_m} \delta_{ct} |\mathscr{I}_c| \hat{T}_{ct}}{\sum_i \delta_{ict}^2} \quad ,$$

where $|\mathscr{I}_c|$ denotes the number of individuals in cohort c. Finally, replacing $\delta_{ct} = \bar{y}_{d(c)t} - \bar{y}_{o(c)t}$ and denoting $w_c \equiv \frac{|\mathscr{I}_c| \cdot \delta_{ct}^2}{\sum_{c \in \mathscr{C}_m} |\mathscr{I}_c| \cdot \delta_{ct}^2}$ allows us to arrive at our desired formula

$$\hat{\mu}_r = \sum_{c \in \mathscr{C}_m} \frac{w_c \hat{T}_{ct}}{\left(\bar{y}_{d(c)t} - \bar{y}_{o(c)t} \right)} \quad ,$$

where $\hat{T}_{ct} \xrightarrow{p} T_{ct}$ implies that $\hat{\mu}_r \xrightarrow{p} \sum_{c \in \mathscr{C}_m} \frac{w_c T_{ct}}{\left(\bar{y}_{d(c)t} - \bar{y}_{o(c)t} \right)}$.

Appendix D: Estimation Details

Description of Estimation Sample

As we describe in Section 5.1, we construct moments using the baseline sample of movers described in Table 1. The resulting enrollee years represent approximately 12,500 cohorts, which are uniquely defined by their origin state, destination state, and move year, and each cohort is observed for an average of 6.5 years. We define the average abuse rate for each cohort and year relative to the move —$\hat{\bar{y}}_{cr}$—and our estimation targets these nearly 80,000 empirical moments, weighting each moment in the estimation procedure by the number of enrollee-years in the sample that were used to construct the moment. We provide summary statistics for these moments in Table A.10.The baseline parameters of our model generate predicted moments for each cohort and year relative to the move, and we describe these calculations in Section 5.1.[17]

Estimation Details

We estimate our model parameters by solving a weighted non-linear least-squares optimization, where the residuals are defined as the weighted distance between the empirical and predicted moments and weights given by the number of patient-years in the sample that we used in the moment's construction. Our solution vector consists of values for the six parameter values which vary by location—$\bar{a}_{j0}$, γ_j, π_j^+, π_j^-, η_j^+, and η_j^-—for each of our fifty-one locations, as well as the separable temporal component of availability τ_t for each of the ten years of our sample.

We estimate our model parameters by solving a weighted non-linear least-squares optimization. Our solution vector consists of values for the six parameter values which vary by location—$\bar{a}_{j0}$, γ_j, π_j^+, π_j^-, η_j^+, and η_j^-—for each of our fifty-one locations, as well as the separable temporal component of availability τ_t for each of the ten years of our sample.

We perform two normalizations in our estimation procedure. First, we fix two sets of values to zero as references: the first calendar year availability parameter and an arbitrary state's person-specific addiction transition effects. We do this in order to avoid collinearity in the sums that are used in the computation of addiction $(\pi_j + \eta_j)$ and availability $(\gamma_j + \tau_t)$ in our model. Thus, the solution vector consists of 313 values.

Second, we set the global mean addict share share in our study population to an arbitrarily chosen $\bar{a} = 0.10$. Specifically, we require our estimation to satisfy

$$\bar{a} = \sum_{m=2007}^{2014} \sum_{c \in \mathscr{C}_m} \hat{\bar{a}}_{cr} \frac{\hat{n}_{cr}}{N},$$

where $\hat{n}_{cr}$ denotes the number of enrollee-years used to construct the moment $\hat{\bar{y}}_{cr}$, $\bar{a}_{cr}$ denotes the average addict share in a cohort and relative year (as calculated in Section 5.1), and N is the total number of enrollee-years in the sample. We perform this normalization because the multiplicative property of our outcome implies that the relative magnitudes of addiction and availability parameters are not meaningful. For example, doubling all addict shares and halving all availability parameters would leave predicted abuse unchanged in our model. Only the geographic variation—the distribution within the addiction and availability—is meaningful. Thus, we arbitrarily fix the magnitude of the global addict shares. In the section below, we show robustness of our parameter estimates to various other global addict shares, and in ongoing work (not reported) we are able to instead include this as an estimation parameter.

[17] Note that in order to avoid specifying the location of cohorts in the year of their move, our baseline estimation proceeds as if the year following the move follows immediately after the year before the move. Thus, cohorts are in their origin states $o(c)$ during relative year -1. Since relative year 1 follows immediately, cohorts are in their destination state $d(c)$ in the following year . At the end of this section we will show robustness to this assumption.

Our actual estimation occurs in three steps. In the first step of our estimation, we set various initial values for the optimization algorithm. Using fifteen sets of initial values, we set all parameter values equal to a constant ranging between 0.00 and 0.15, including the end-points and evenly distributed between. Second, we solve the weighted least squares problem using the Levenberg-Marquardt algorithm separately for each of these fifteen initial values (Moré 1978). Taking the penalty approach to non-linear optimization, we enforce two constraints: (i) that all addiction parameters and availability parameters that an individual is subject to in the model are bounded between 0 and 1, and (ii) that the global mean addict share is equal to $\bar{a} = 0.10$. 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.

Robustness

In this section, we examine several extensions and modifications of our model and show in Table A.11 that our parameter estimates across these extensions are highly correlated with estimates from our baseline model. For example, in column (1), we examine the model specification which includes the year of move in the computation of addict shares among move cohorts. Thus, mover do not move instantaneously from their origins in relative year -1 to the destinations in relative year 1, but instead undergo another year of transitioning to and from addiction. In this extension, we assume that the addict shares of move cohorts evolve according to the place-specific addiction transition parameters of their origin during their move year.

In column (2), we allow for movers to differ from non-movers through relative year fixed effects, which we denote by ρ_r. We allow these to enter through the availability channel. Thus, opioid abuse among addicts in this specification simply becomes

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

In column (3), we examine the specification where we allow for non-addicts to abuse opioids at a constant rate across time and geographies. Denoting this constant rate as λ, we have that the average abuse rate for cohort $c$in time t becomes

$$\bar{y}_{j(c,t)t} = \bar{a}_{ct} \cdot \gamma_{j(c,t)t} + (1 - \bar{a}_{ct}) \cdot \lambda.$$

We estimate this parameter λ as an additional parameter in our model estimation. This yields an estimate of λ,the probability of abuse among non-addicts, to be 1.2%. As a point of comparison, the average probability of abuse among addicts—computing $\bar{\gamma}_{jt}$ by averaging equally across calendar years and states—is 34.3%.

Finally, in columns (4) and (5), we consider modifications of our baseline model which use various specifications of the global addict share $\bar{a}$.

Appendix E: Appendix Figures and Tables

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

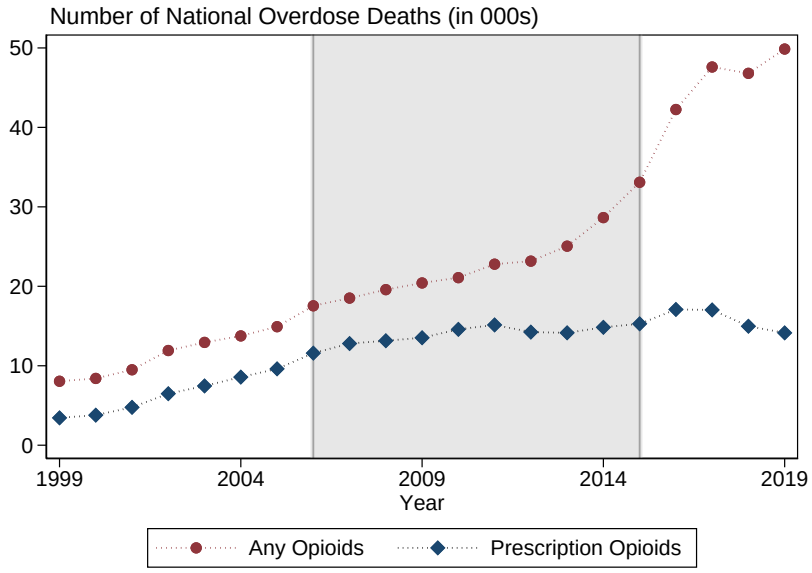

Notes: Figure presents the number of national opioid overdose deaths and the subset of those which involved prescription opioids. The years of our sample – 2006 through 2015 – are shaded. The data was directly downloaded from the CDC WONDER Multiple Cause of Death (MCD) Files (NCHS 2021).

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

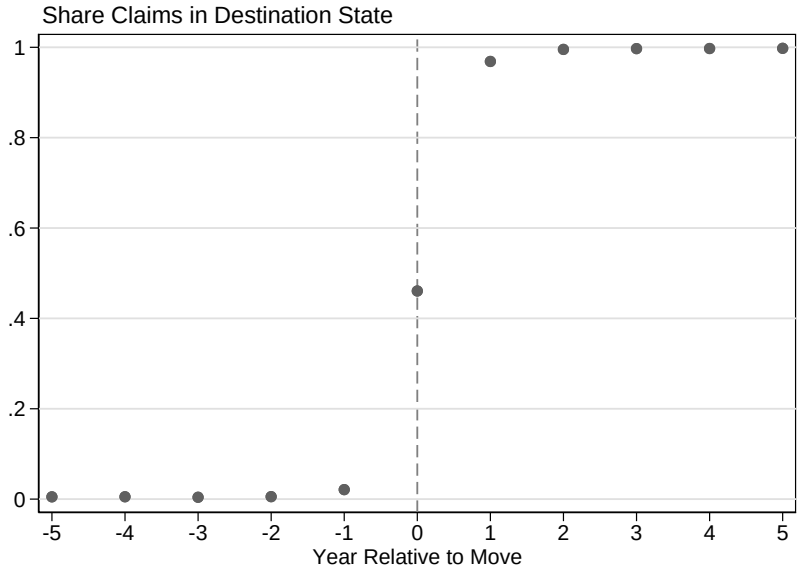

Notes: Figure presents the average fraction of claims in a mover's destination state (as a fraction of claims in either the origin or destination state) by year relative to move. Observations are at the mover-year level. The figure shows a sharp change in the year of the move, with only a small share 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: Event Studies - Other Outcomes

(a) Baseline Measure

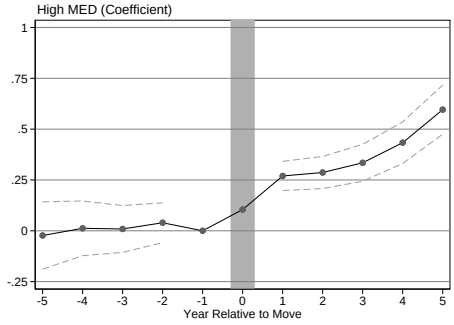

(b) Many Prescribers

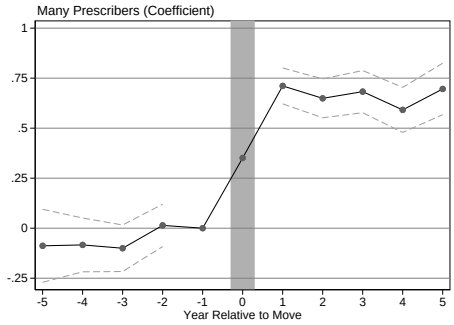

(c) Overlapping Prescriptions

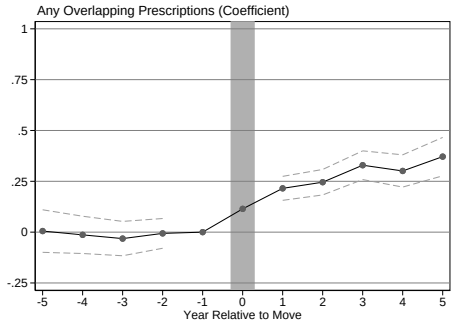

Notes: Figure shows the weighted average across move years of estimates from equation (3) of $\mu_{r(i,t)}$—the coefficient on year relative to move interacted with $\hat{\delta}_{ct}$—where the coefficient for relative year -1 is normalized to 0 and weights are given by sample size. The dependent variable $\tilde{y}_{ict}$ is the residual of the mover's abuse rate with the indicated opioid use measure after subtracting the predicted abuse rate using origin-year and five-year age bin fixed effects as estimated from non-movers. The dashed lines are upper and lower bounds of the 95% confidence interval. We construct this confidence interval following a three-step procedure. In the first step, for each state-year, we construct the asymptotic distribution of its average rate of abuse $\bar{y}_{jt}$ according to the central limit theorem, which is a normal distribution with mean and standard deviation calculated from the data. In the second step, we bootstrap equation (3) with 50 repetitions drawn at the enrollee-year level, making a random draw from the asymptotic distribution for each mover's origin, destination, and year to construct their $\hat{\delta}_{ct}$ for each repetition. Finally, using these bootstraps to estimate the standard errors for each move year estimate, we calculate the variance of the weighted average coefficients by assuming that the move year coefficients are independent. 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 abuse measures of Many Prescribers and Overlapping Prescriptions. These abuse measures are discussed in Appendix A and sample sizes are given in Table A.1.

Figure A.4: Event Studies - Robustness

(a) Baseline

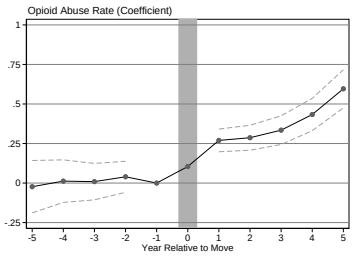

(b) Commuting Zone

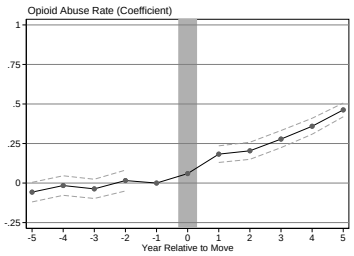

(c) County

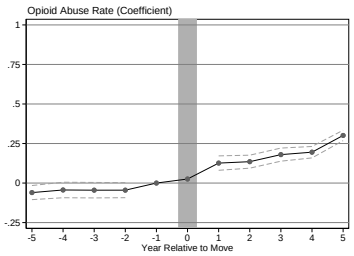

(d) Individual FE

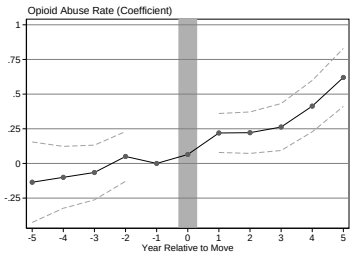

(e) Relative Year FE

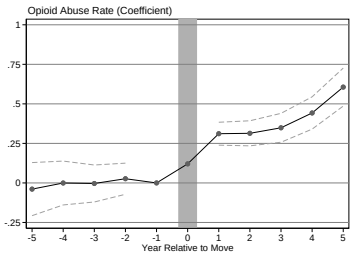

(f) Individual and Relative Year FE

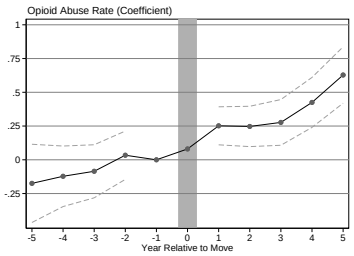

(g) Part D Full Coverage

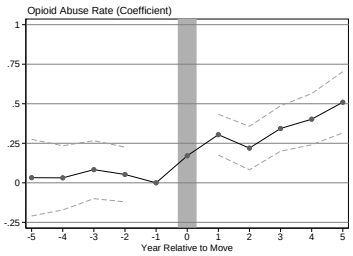

(h) Alive All Years

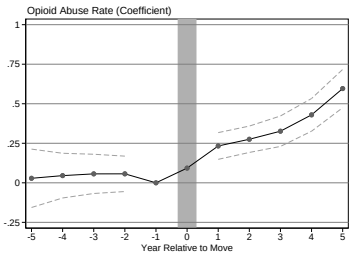

Notes: Figure shows the weighted average across move years of estimates from equation (3) of $\mu_{r(i,t)}$—the coefficient on year relative to move interacted with $\hat{\delta}_{ct}$—where the coefficient for relative year -1 is normalized to 0 and weights are given by sample size. The dependent variable $\tilde{y}_{ict}$ is the residual of the mover's abuse rate after subtracting the predicted abuse rate using origin-year and five-year age bin fixed effects as estimated from non-movers. The dashed lines are upper and lower bounds of the 95% confidence interval. We construct this confidence interval following a three-step procedure. In the first step, for each state-year, we construct the asymptotic distribution of its average rate of abuse $\bar{y}_{jt}$ according to the central limit theorem, which is a normal distribution with mean and standard deviation calculated from the data. In the second step, we bootstrap equation (3) with 50 repetitions drawn at the enrollee-year level, making a random draw from the asymptotic distribution for each mover's origin, destination, and year to construct their $\hat{\delta}_{ct}$ for each repetition. Finally, using these bootstraps to estimate the standard errors for each move year estimate, we calculate the variance of the weighted average coefficients by assuming that the move year coefficients are independent. Panel A is our baseline specification, panel B uses commuting zones as the moving geography, and panel C uses counties as the moving geography. Panel D uses the specification in equation (3) but includes individual fixed effects, panel E instead includes relative year fixed effects, and panel F includes both individual and relative year fixed effects. Panel G restricts to movers with 12 months of Part D coverage for each year they are observed in the sample, and panel H restricts to all movers who did not die from 2006 to 2015. These specifications are discussed in Appendix B, and sample sizes are given in Table A.7.

Figure A.5: Event Studies - Selection and Attrition

(a) Part D Extensive Margin

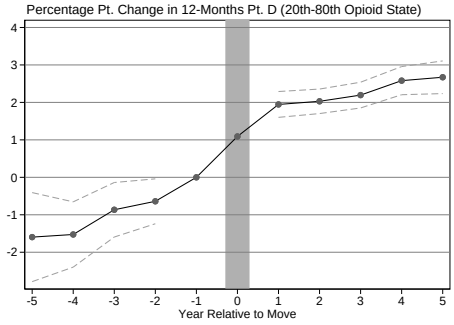

(b) Died

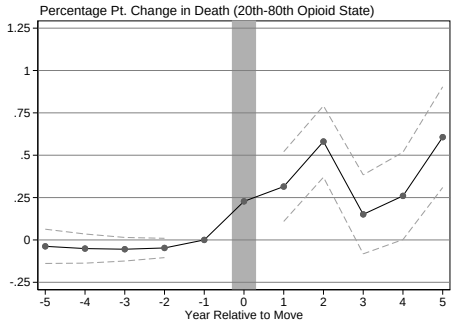

Notes: Figure shows the weighted average across move years of estimates from equation (3) of $\mu_{r(i,t)}$—the coefficient on year relative to move interacted with $\hat{\delta}_{ct}$—where the coefficient for relative year -1 is normalized to 0 and weights are given by sample size. The dependent variable $\tilde{y}_{ict}$ is the residual of the mover's outcome after subtracting the predicted outcome using origin-year and five-year age bin fixed effects as estimated from non-movers, where this outcome as having twelve complete months of Part D coverage during the year (panel A) and an indicator for death during the year (panel B). The dashed lines are upper and lower bounds of the 95% confidence interval. We construct this confidence interval following a three-step procedure. In the first step, for each state-year, we construct the asymptotic distribution of its average rate of abuse $\bar{y}_{jt}$ according to the central limit theorem, which is a normal distribution with mean and standard deviation calculated from the data. In the second step, we bootstrap equation (3) with 50 repetitions drawn at the enrollee-year level, making a random draw from the asymptotic distribution for each mover's origin, destination, and year to construct their $\hat{\delta}_{ct}$ for each repetition. Finally, using these bootstraps to estimate the standard errors for each move year estimate, we calculate the variance of the weighted average coefficients by assuming that the move year coefficients are independent. The y-axis is scaled to represent a 20th percentile to 80th percentile move in the rate of opioid abuse. The sample in panel A is all movers, in addition to those without 12 complete months of Part D coverage (N = 662,243 enrollee-years). The sample in panel B is all movers, in addition to the year of death for those who die in the sample (N = 533,185 enrollee-years).

Figure A.6: Time Trends in Opioid Prescribing Rates

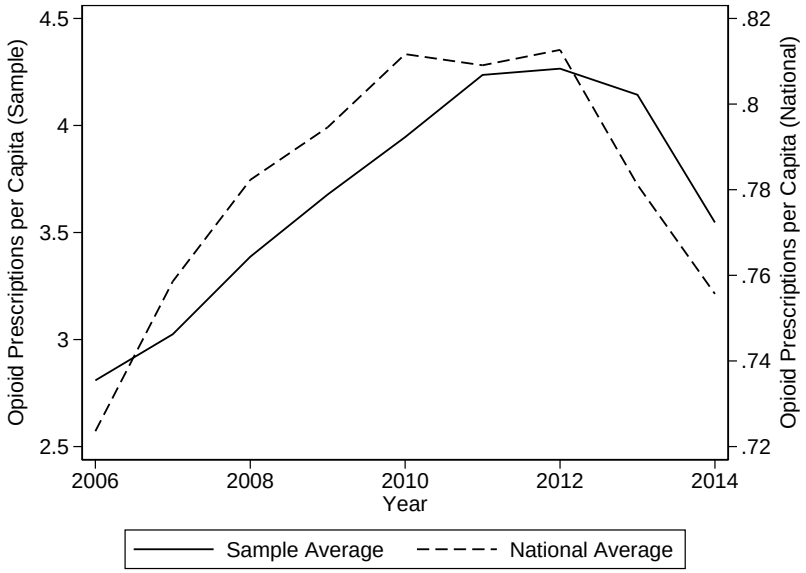

Notes: Figure presents averages over time for 2006-2014 of the number of prescriptions per capita among non-movers and from national QuintilesIMS data. Both prescription rates are calculated on a January-December calendar year.

Figure A.7: Event Study – Poisoning Event

(a) All Users

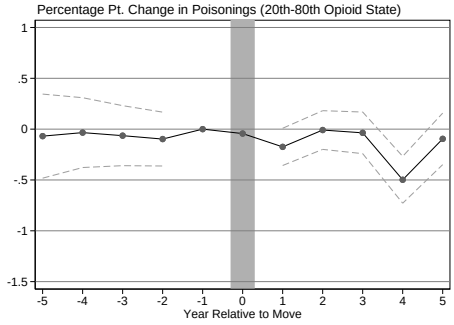

(b) Naive Users

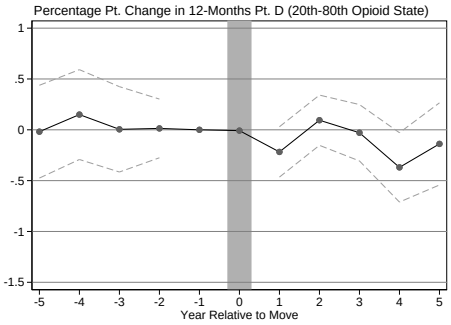

(c) Prior Users

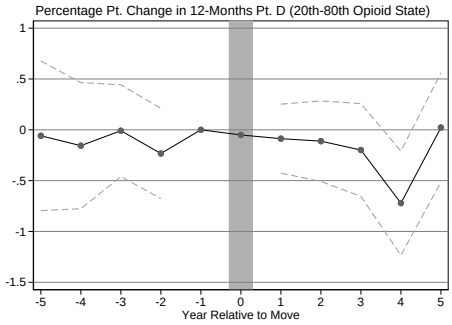

Notes: Figure shows the weighted average across move years of estimates from equation (3) of $\mu_{r(i,t)}$—the coefficient on year relative to move interacted with $\hat{\delta}_{ct}$—where the coefficient for relative year -1 is normalized to 0 and weights are given by sample size. The dependent variable $\tilde{y}_{ict}$ is the residual of the mover's outcomes after subtracting the predicted abuse rate using origin-year and five-year age bin fixed effects as estimated from non-movers, where the outcome is defined as an indicator for whether an individual experienced a non-fatal poisoning event in that year, defined as an ER visit or inpatient admission with an ICD-9 diagnosis code indicating a poisoning. Values on the y-axis refer to the percentage point change in a poisoning event associated with a move from the 20th to 80th percentile state with respect to the rate of opioid abuse. The dashed lines are upper and lower bounds of the 95% confidence interval. We construct this confidence interval following a three-step procedure. In the first step, for each state-year, we construct the asymptotic distribution of its average rate of abuse $\bar{y}_{jt}$ according to the central limit theorem, which is a normal distribution with mean and standard deviation calculated from the data. In the second step, we bootstrap equation (3) with 50 repetitions drawn at the enrollee-year level, making a random draw from the asymptotic distribution for each mover's origin, destination, and year to construct their $\hat{\delta}_{ct}$ for each repetition. Finally, using these bootstraps to estimate the standard errors for each move year estimate, we calculate the variance of the weighted average coefficients by assuming that the move year coefficients are independent. Sample sizes are given in Table A.9.

Figure A.8: Model Calibrations

(a) Movers

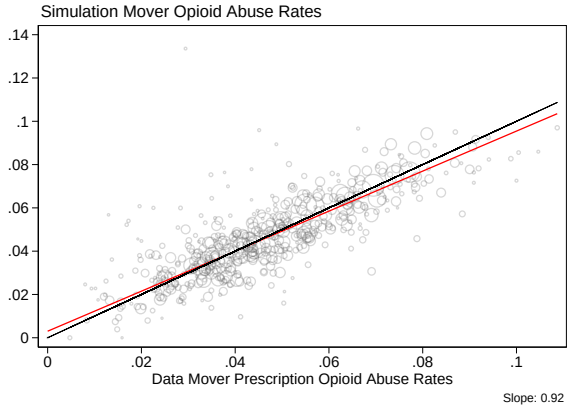

(b) Non-movers

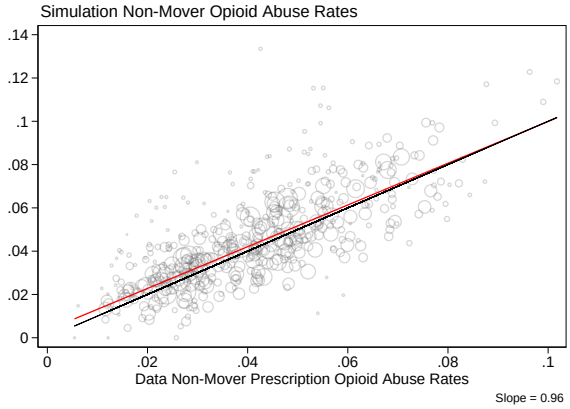

Notes: These figures compare mover (panel A) and non-mover (panel B) mean opioid abuse rates predicted from our simulations to those in the Medicare data. Each moment is a (state, calendar year) combination. In each panel, the y-axis corresponds to the mean of each moment across 50 simulations, and the black 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 red 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. Notes for each figure show the slope of the line of best fit.

Figure A.9: Time Series of Average Opioid Abuse Rates

(a) Medicare Data

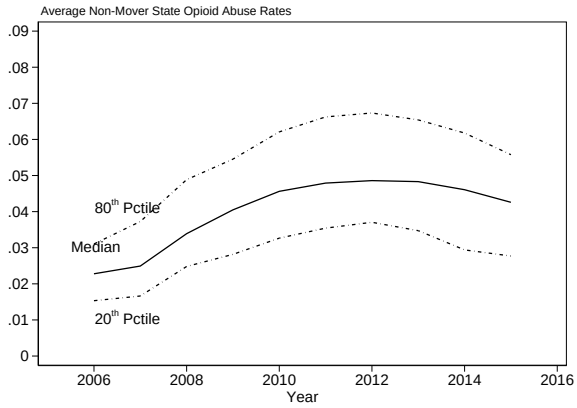

(b) Model Predictions

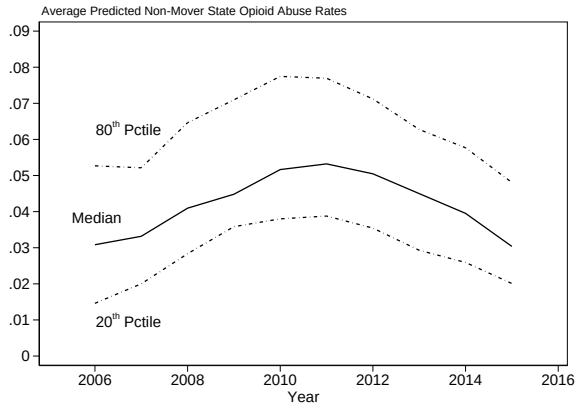

Notes: These figures show time series of average abuse rates among non-movers in the 20th percentile, 50th percentile, and 80th percentile states in each year. Percentiles are calculated for each year, and each line corresponds to different states in each year. Panel A shows results from the Medicare data, while panel B shows predicted non-mover moments from parameters estimated in our state-level movers only model. Opioid abuse is measured using our baseline measure, as defined inAppendix A.

Table A.1: Event Study Coefficients – Other Outcomes

	(1)	(2)	(3)
	All	Naive	Prior User
Panel A: Baseline Measure			
Average of 1-5 years post-move	0.384	0.156	0.667
	(0.047)	(0.032)	(0.112)
Enrollees	90,890	32,648	30,460
Enrollee-years	521,523	219,100	202,309
Panel B: Many Prescribers			
Average of 1-5 years post-move	0.667	0.340	1.204
	(0.054)	(0.056)	(0.119)
Enrollees	90,890	32,648	30,460
Enrollee-years	521,523	219,100	202,309
Panel C: Overlapping Prescriptions			
Average of 1-5 years post-move	0.293	0.240	0.458
	(0.037)	(0.038)	(0.080)
Enrollees	90,890	32,648	30,460
Enrollee-years	521,523	219,100	202,309

Notes: Table reports the average of coefficients and bootstrapped standard errors (in parentheses) in relative years 1-5 for the sample of all movers, naive enrollees, and prior users. Estimates correspond to the weighted average across move years of estimates from equation (3) of $\mu_{r(i,t)}$—the coefficient on year relative to move interacted with $\hat{\delta}_{ct}$—where the coefficient for relative year -1 is normalized to 0 and weights are given by sample size. Each panel reports results for a different outcome: our baseline measure (high morphine equivalent doses or "High MED"), Many Prescribers, and Overlapping Prescriptions. These opioid abuse measures are discussed in Appendix A. The omitted category is relative year -1. 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. Except in the sample of all movers, we omit the approximately 20% of enrollee-years for enrollees who are not observed in the year before moving. Standard errors are calculated following a three-step procedure. In the first step, for each state-year, we construct the asymptotic distribution of its average rate of abuse $\bar{y}_{jt}$ according to the central limit theorem, which is a normal distribution with mean and standard deviation calculated from the data. In the second step, we bootstrap equation (3) with 50 repetitions drawn at the enrollee-year level, making a random draw from the asymptotic distribution for each mover's origin, destination, and year to construct the $\hat{\delta}_{ct}$ for each repetition. Finally, using these bootstraps to estimate the standard errors for each move year estimate, we calculate the standard deviation of the weighted average coefficients by assuming that the move year coefficients are independent.

Table A.2: Event Study Coefficients - Balanced Panels

	Observed Post-Move for at Least					
	Baseline	1 Year	2 Years	3 Years	4 Years	5 Years
1 year post-move	0.270	0.273	0.270	0.264	0.243	0.270
	(0.037)	(0.038)	(0.036)	(0.042)	(0.052)	(0.057)
5 years post-move	0.596	0.599	0.617	0.616	0.586	0.545
	(0.062)	(0.031)	(0.029)	(0.036)	(0.048)	(0.054)
Enrollees	90,890	87,607	79,025	61,132	46,020	33,394
Enrollee-years	521,523	510,931	486,347	405,380	325,267	245,980

Notes: Table reports the average of coefficients and bootstrapped standard errors (in parentheses) in relative years 1 and 5 for our baseline sample as well as subsamples where units are observed in the post-move period for at least the number of years indicated. Estimates correspond to the weighted average across move years of estimates from equation (3) of $\mu_{r(i,t)}$—the coefficient on year relative to move interacted with $\hat{\delta}_{ct}$—where the coefficient for relative year -1 is normalized to 0 and weights are given by sample size. This balancing follows Callaway and Sant'Anna (2020). Standard errors are calculated following a three-step procedure. In the first step, for each state-year, we construct the asymptotic distribution of its average rate of abuse $\bar{y}_{jt}$ according to the central limit theorem, which is a normal distribution with mean and standard deviation calculated from the data. In the second step, we bootstrap equation (3) with 50 repetitions drawn at the enrollee-year level, making a random draw from the asymptotic distribution for each mover's origin, destination, and year to construct the $\hat{\delta}_{ct}$ for each repetition. Finally, using these bootstraps to estimate the standard errors for each move year estimate, we calculate the standard deviation of the weighted average coefficients by assuming that the move year coefficients are independent.

Table A.3: State-Level Opioid Abuse Measure Summary Statistics

	25th Percentile	Median	75th Percentile	Interquartile Range
Baseline Measure	0.028	0.040	0.054	0.026
Many Prescribers	0.044	0.058	0.073	0.030
Overlapping Prescriptions	0.114	0.146	0.175	0.061

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

Table A.4: 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: This table identifies opioid conversion factors used in construction of our baseline measure of abuse. Adapted from Meara et al. (2016) supplementary material as well as http://www.agencymeddirectors.wa.gov/calculator/dosecalculator.htm.

Table A.5: Measures of Hazardous Opioid Prescriptions

	Baseline Measure	Many Prescribers	Overlapping Prescriptions	Any Opioid	Chronic Opioid Use
Baseline Measure	1.000				
Many Prescribers	0.714	1.000			
Overlapping Prescriptions	0.729	0.797	1.000		
Any Opioid	0.486	0.690	0.929	1.000	
Chronic Opioid Use	0.533	0.659	0.948	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 Opioid, and Chronic Opioid Use variables for 2006-2015, averaged across years at the state level. These various opioid abuse measures are discussed in Appendix A. State-level correlations are weighted by the 2010 Census population.

Table A.6: Opioid Prescribing Measures and Adverse Opioid Outcomes

(a) Opioid Prescribing Measures

	All Enrollees	Enrollees with Poisoning Event in Year $t+1$	Enrollees with Opioid Use Disorder in Year $t+1$
Baseline Measure	0.044	0.335	0.282
Many Prescribers	0.061	0.379	0.338
Overlapping Prescribers	0.154	0.676	0.567
Any Opioids	0.420	0.905	0.808
Chronic Opioid Use	0.210	0.733	0.610
Share of Non-Movers	1.000	0.002	0.011
N	4,678,601	7,195	49,399

(b) Predictive Power: Positive Likelihood Ratios

	Poisoning Event in Year $t+1$	Opioid Use Disorder in Year $t+1$
Baseline Measure	7.67	7.12
Many Prescribers	5.76	5.58
Overlapping Prescribers	4.37	3.88
Any Opioids	2.12	1.95
Chronic Opioid Use	3.54	3.12

Notes: These tables show the relationship between adverse opioid outcomes (opioid poisonings and diagnoses of opioid use disorders) with our proxies for opioid abuse and opioid prescribing measures. Panel A shows the rate of these measures among non-mover enrollee years covered by traditional Medicare for the full year (for which we observe the following year). The first column shows all enrollee-years while the second and third columns restrict to enrollees with adverse opioid outcomes in the year following. Panel B shows the positive likelihood ratios defined for each opioid measure and adverse outcome. The positive likelihood ratio is defined as the ratio of true to false positives when the opioid abuse or prescribing measure is used as a diagnostic for adverse opioid outcomes in the following year. These various opioid abuse measures are discussed in Appendix A.

Table A.7: Event Study Coefficients – Robustness

	(1)	(2)	(3)
	All	Naive	Prior User
Panel A: Baseline Geography (State)			
Average of 1-5 years post-move	0.384	0.155	0.667
	(0.047)	(0.032)	(0.112)
Enrollees	90,890	32,648	30,460
Enrollee-years	521,523	219,100	202,309
Panel B: Commuting Zone			
Average of 1-5 years post-move	0.297	0.070	0.494
	(0.026)	(0.020)	(0.080)
Enrollees	120,760	45,001	39,434
Enrollee-years	694,963	303,379	260,660
Panel C: County			
Average of 1-5 years post-move	0.189	0.0401	0.298
	(0.020)	(0.016)	(0.060)
Enrollees	131,790	49,958	41,788
Enrollee-years	755,233	336,787	273,475
Panel D: Individual FE			
Average of 1-5 years post-move	0.348	0.144	0.584
	(0.087)	(0.053)	(0.180)
Enrollees	90,890	32,648	30,460
Enrollee-years	521,523	219,100	202,309
Panel E: Relative Year FE			
Average of 1-5 years post-move	0.405	0.143	0.692
	(0.048)	(0.031)	(0.111)
Enrollees	90,890	32,648	30,460
Enrollee-years	521,523	219,100	202,309
Panel F: Individual and Relative Year FE			
Average of 1-5 years post-move	0.366	0.128	0.635
	(0.087)	(0.053)	(0.181)
Enrollees	90,890	32,648	30,460
Enrollee-years	521,523	219,100	202,309
Panel G: Part D All Years			
Average of 1-5 years post-move	0.356	0.161	0.603
	(0.078)	(0.043)	(0.157)
Enrollees	20,869	11,123	9,746
Enrollee-years	168,179	89,812	78,367
Panel H: Alive All Years			
Average of 1-5 years post-move	0.372	0.147	0.638
	(0.050)	(0.033)	(0.1105)
Enrollees	76,332	27,434	24,779
Enrollee-years	453,963	191,904	172,834

Notes: Table reports the average of coefficients and bootstrapped standard errors (in parentheses) in relative years 1-5 for the sample of all movers, naive enrollees, and prior users. Estimates correspond to the weighted average across move years of estimates from equation (3) of $\mu_{r(i,t)}$—the coefficient on year relative to move interacted with $\hat{\delta}_{ct}$—where the coefficient for relative year -1 is normalized to 0 and weights are given by sample size. 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. Panel D uses the specification in equation (3) but includes individual fixed effects, panel E instead includes relative year fixed effects, and panel F includes both individual and relative year fixed effects. Panel G restricts to movers with 12 months of Part D coverage for each year they are observed in the sample, and panel H restricts to all movers who did not die from 2006 to 2015. There are about 700 commuting zones and 3,000 counties in the sample. The omitted category is relative year -1. 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. Except in the baseline sample, we omit the approximately 20% of enrollee-years for enrollees who are not observed in the year before moving. Standard errors are calculated following a three-step procedure. In the first step, for each state-year, we construct the asymptotic distribution of its average rate of abuse $\bar{y}_{jt}$ according to the central limit theorem, which is a normal distribution with mean and standard deviation calculated from the data. In the second step, we bootstrap equation (3) with 50 repetitions drawn at the enrollee-year level, making a random draw from the asymptotic distribution for each mover's origin, destination, and year to construct the $\hat{\delta}_{ct}$ for each repetition. Finally, using these bootstraps to estimate the standard errors for each move year estimate, we calculate the standard deviation of the weighted average coefficients by assuming that the move year coefficients are independent.

Table A.8: Event Study Coefficients – Moves Up and Moves Down

	(1)	(2)	(3)
	All	Naive	Prior User
Panel A: Moves Up			
1 year post-move	0.045	0.109	-0.170
	(0.048)	(0.027)	(0.125)
5 years post-move	0.638	0.395	1.111
	(0.084)	(0.076)	(0.223)
Enrollees	49,857	18,243	16,212
Enrollee-years	284,929	122,464	107,671
Panel B: Moves Down			
1 year post-move	0.589	-0.018	1.213
	(0.056)	(0.029)	(0.129)
5 years post-move	0.584	0.119	1.073
	(0.082)	(0.053)	(0.210)
Enrollees	41,033	14,405	14,248
Enrollee-years	236,594	96,636	94,638

Notes: Table reports the average of coefficients and bootstrapped standard errors (in parentheses) in relative years 1 and 5 for the sample of all movers, naive enrollees, and prior users. Estimates correspond to the weighted average across move years of estimates from equation (3) of $\mu_{r(i,t)}$—the coefficient on year relative to move interacted with $\hat{\delta}_{ct}$—where the coefficient for relative year -1 is normalized to 0 and weights are given by sample size. The table reports estimates for moves up and moves down based on the difference in the rates of opioid abuse between the origin and destination state in the year of the move. The omitted category is relative year -1. 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. Except in Column (1), we omit the approximately 20% of enrollee-years for enrollees who are not observed in the year before moving. Standard errors are calculated following a three-step procedure. In the first step, for each state-year, we construct the asymptotic distribution of its average rate of abuse $\bar{y}_{jt}$ according to the central limit theorem, which is a normal distribution with mean and standard deviation calculated from the data. In the second step, we bootstrap equation (3) with 50 repetitions drawn at the enrollee-year level, making a random draw from the asymptotic distribution for each mover's origin, destination, and year to construct the $\hat{\delta}_{ct}$ for each repetition. Finally, using these bootstraps to estimate the standard errors for each move year estimate, we calculate the standard deviation of the weighted average coefficients by assuming that the move year coefficients are independent.

Table A.9: Event Study - Poisoning Event

	(1)	(2)	(3)
	All	Naive	Prior User
Average of 1-5 years post-move	-0.163	-0.132	-0.219
	(0.108)	(0.155)	(0.229)
Enrollees	90,890	32,648	30,460
Enrollee-years	521,523	219,100	202,309

Notes: Table reports the average of coefficients and bootstrapped standard errors (in parentheses) in relative years 1-5 for the sample of all movers, naive enrollees, and prior users. Estimates correspond to the weighted average across move years of estimates from equation (3) of $\mu_{r(i,t)}$—the coefficient on year relative to move interacted with $\hat{\delta}_{ct}$—where the coefficient for relative year -1 is normalized to 0 and weights are given by sample size. The outcome is an indicator for whether an individual experienced any non-fatal poisoning event, defined as an ER visit or inpatient admission with an ICD-9 diagnosis code indicating a poisoning. Coefficient is scaled to represent a 20th to 80th percentile move in the opioid abuse rate. The omitted category is relative year -1. 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. Except in the baseline sample, we omit the approximately 20% of enrollee-years for enrollees who are not observed in the year before moving. Standard errors are calculated following a three-step procedure. In the first step, for each state-year, we construct the asymptotic distribution of its average rate of abuse $\bar{y}_{jt}$ according to the central limit theorem, which is a normal distribution with mean and standard deviation calculated from the data. In the second step, we bootstrap equation (3) with 50 repetitions drawn at the enrollee-year level, making a random draw from the asymptotic distribution for each mover's origin, destination, and year to construct the $\hat{\delta}_{ct}$ for each repetition. Finally, using these bootstraps to estimate the standard errors for each move year estimate, we calculate the standard deviation of the weighted average coefficients by assuming that the move year coefficients are independent.

Table A.10: Estimation Sample Summary Statistics

Sample Used to Construct Moments	
Number of Movers	90,301
Number of Cohorts	12,528
Number of Enrollee-Years	451,124
Average Abuse Rate	4.8%
Moments Targeted in Model Estimation (Years Relative to Move for Each Cohort)	
Number of Moments	81,892
Median Number of Enrollee-Years per Moment	2
Mean Number of Enrollee-Years per Moment	5.5
Minimum Number of Enrollee-Years per Moment	1
Maximum Number of Enrollee-Years per Moment	316

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 Table 1 column 1, excluding enrollee-years observed during the year of move. The empirical moments are average abuse rates constructed at the cohort – relative-year level, where cohorts are defined by a mover's origin, destination, and move year.

Table A.11: Robustness of Parameter Estimates to Model Extensions

	(1) Allowing Addiction Transitions During Relative Year 0	(2) Including Relative Year Fixed Effects	(3) Allowing for Abuse by Non-Addicts	(4) Alternative Global Addict Share ($\bar{a} = 0.15$)	(5) Alternative Global Addict Share ($\bar{a} = 0.20$)
	Correlation Coefficient with Baseline Model Parameter Estimates				
a_{j0}	0.95	0.97	0.91	0.95	0.78
γ_j	0.75	0.98	0.94	0.97	0.82
π_j^+	0.87	0.95	0.94	0.89	0.84
π_j^-	0.83	0.91	0.92	0.92	0.84
η_j^+	0.85	0.81	0.87	0.66	0.49
η_j^-	0.90	0.88	0.84	0.81	0.76
τ_t	0.89	1.00	1.00	1.00	0.96

Notes: Table presents correlation coefficients between the baseline model parameter estimates and various model extensions and modifications. Column (1) examines parameter estimates derived from allowing for addiction transitions to occur for an additional year in the year of the move. These transitions are assumed to occur according to the place-specific addiction transition parameters of the origin location. Column (2) examines parameter estimates derived from estimating a version of the model which allows for movers to differ from non-movers through relative year fixed effects, which enter through the availability channel. Column (3) allows for abuse by non-addicts at a constant rate, which is an additional parameter that we estimate. Finally, columns (4) and (5) specify different global addict shares ($\bar{a} = 0.15$ and $\bar{a} = 0.20$). We discuss these specifications further in Appendix D.