

OPIOIDS FOR THE MASSES: WELFARE TRADEOFFS IN THE REGULATION OF NARCOTIC PAIN MEDICATIONS*

Angela E. Kilby

December 2021

Abstract

Use of prescription opioid pain relievers to manage pain increased threefold from 2001-2011, as medical guidelines were revised to emphasize that appropriate pain management is required for an acceptable standard of care. However, a concomitant rapid rise in opioid abuse, addiction, overdose, and death has led to escalating policy efforts to crack down on opioid prescribing. This paper sheds light on the tradeoffs of public policies that reduce the supply of medical opioids by investigating their health, labor, and welfare ramifications. I exploit state-level variation in the introduction of Prescription Drug Monitoring Program (PDMP) laws and make use of several rich data sources, documenting that PDMPs reduce the distribution of opioids, and achieve key policy goals by reducing rates of opioid use disorder, and reducing opioid overdose deaths by about 8%. I also leverage rich individual-level medical claims data and work records to document substantial costs resulting from these policies, including more missed days of work for injured and disabled individuals, and substitution towards more-expensive medical care during medical episodes requiring pain management. A rough back-of-the-envelope welfare calculation suggests the welfare losses and gains from regulation are on the same order of magnitude – approximately \$8.2–\$10.9 billion per year in increased costs from inpatient and outpatient medical spending plus lost wages, compared to \$4.8–\$7.6 billion per year in benefits from lives saved from opioid overdose.

*Department of Economics, Northeastern University, 360 Huntington Avenue, 301 Lake Hall, Boston, MA 02115. Email: a.kilby@northeastern.edu. I am grateful to Jon Gruber, Amy Finkelstein, and Heidi Williams for their advice, guidance and support throughout this project. I also thank Alex Bartik, Enrico Cantoni, Joseph Doyle, Ludovica Gazzè, Sally Hudson, Manisha Padi, Jonathan Petkun, Jim Poterba, Brendan Price, Otis Reid, Natalia Rigol, Adam Sacarny, Ashish Shenoy, Evann Smith, Melanie Wasserman, and the MIT Public Finance and Summer Applied Micro Lunch participants, for their helpful feedback. Funding support is gratefully acknowledged from the National Science Foundation Graduate Research Fellowship, the National Institute on Aging Pre-Doctoral Fellowship, and the National Institute on Aging Post-Doctoral Fellowship.

1 Introduction

Prescription opioid pain relievers are a highly effective medical technology for the general relief of pain, but they also can be addictive, and their use may lead to dependence, abuse, overdose, and death. Until the mid-1990s, prescription opioids were utilized primarily for acute pain and pain related to terminal cancer, but rarely for chronic, non-cancer pain. Fears of abuse meant that even for acute pain and cancer patients, pain was often undertreated. In the face of concerns that undertreatment of pain was a “serious public health issue” (Federation, 1998), medically indicated use of these drugs over the past 20 years has increased dramatically, and attitudes have liberalized towards the use of opioids for chronic non-cancer pain.

However, in recent years, there has been a considerable backlash against the use of opioids, as a concomitant, and largely unanticipated, rise in adverse events associated with opioid pain relievers has raised alarm in medical and public health communities (Kolodny et al., 2015).¹ As can be seen in Figure 1, a threefold increase from 2001 to 2011 in Morphine Milligram Equivalent (MME) of opioids distributed coincided with a threefold increase in drug overdose deaths linked to opioids during that same time period. Opioid overdose deaths have only continued to rise, and overtook automobile accidents as a leading cause of death in 2016. Policymakers including the Secretary of Health and Human Services have declared rising rates of prescription opioid drug abuse an ‘epidemic’ (US Department of Health and Human Services, 2015). These intensifying concerns prompted the Centers for Disease Control (CDC) to release opioid prescribing guidelines in March 2016, and these guidelines endorsed a major shift in practice, recommending a sharp reduction in opioid prescribing for both acute and chronic pain compared to what was common at that point in time. (Frieden and Houry, 2016).²

In this paper I exploit state-level variation in the timing of one major category of efforts to crack down on the distribution of opioids inside the medical care system – the implementation of statewide Prescription Drug Monitoring Programs (PMPs or PDMPs) – to identify the causal effects of a reduction in opioid supply on health outcomes, health spending, work output, and death, and attempt to integrate these results in order to characterize the welfare tradeoffs inherent in regulating these controlled substances.

PDMPs are state-run databases used to track prescribing and dispensing of controlled prescrip-

¹Chronic pain researchers who were influential in the adoption of opioids for the treatment of chronic pain (Portenoy and Foley, 1986) have later revised their views (Catan and Perez, 2012).

² Under the guidelines, opioids are not preferred for chronic pain, and there are dosage recommendations if utilized of 50 MME per day (and at most 90 MME per day); opioids for acute pain should be limited to three days in most cases, and at most seven days. Prior to these guidelines, it was common for opioid prescriptions to exceed these amounts. For example, in the Marketscan medical claims data used in this paper, 56% of all acute opioid prescriptions exceeded the three-day recommendation, and 15% of all acute prescriptions exceeded seven days; 19% of all chronic prescriptions exceeded the 50 MME per day recommendation, and 9% of all chronic prescriptions exceeded 90 MME per day.

tion drugs to patients, and the data collected are accessible by physicians, pharmacists, and sometimes law enforcement officials. By monitoring the behavior of both doctors and patients, PDMPs aim to reduce undesirable doctor behavior (ranging from overprescribing beyond recommended guidelines, to extreme cases of “pill mills,” where opioids are prescribed to patients without medical need), as well as undesirable patient behavior (i.e., “doctor-shopping”). They are intended to increase physician caution, improve clinical decision-making, and reduce the overall propensity of doctors to prescribe and refill opioids for pain management (PDMP Center of Excellence, 2014). PDMPs are considered to be one of the first policy levers for states to employ to reduce opioid-related harms; in 2002, the DOJ started a grant program to states to support the implementation of state Prescription Drug Monitoring Programs, and they were recommended in the National Drug Control Strategy starting in the mid-2000s (White House, 2005). The March 2016 CDC opioid prescribing guidelines recommended that clinicians always consult their PDMP before initiating opioid therapy and periodically during long-term chronic therapy (Centers for Disease Control and Prevention, 2015).

However, in recent years, efforts to restrict opioid prescribing, ranging from enforcing increasingly rigid PDMP clinician usage rules, to prescription dosage limits set by state legislatures, to the new CDC prescribing guidelines, have engendered some concern and controversy. For example, in a comment on the initial draft, the American Medical Association stated that the new CDC guidelines and supporting discussion were “devoid of a patient-centered view and any real acknowledgement or empathy of the problems chronic pain patients may face” (Ault, 2015). In 2019, acknowledging the substantial impact the guidelines have had on restricting access to opioid pain relievers and some unintended consequences, the CDC released an advisory statement warning about heavy-handed application of its 2016 guidelines, clarifying they do not support hard dosage limits, or sudden discontinuation/“cutting off” of patients. (Centers for Disease Control, 2019a). The American Medical Association recently argued that the CDC guidelines have increased stigma for pain patients and resulted in legitimate pain care being denied, and as such should be substantially revised (Madara, 2020).

Underlying this controversy over optimal policy are complex welfare considerations. Early understanding of the ‘opioid epidemic’ conceptualized it as resulting from diversion and theft of opioids from the medical care system for the purpose of illicit abuse, and emphasized that legitimate medical use of opioids for pain is safe and effective. A welfare calculation under this scenario would entail valuing costs and benefits of increased regulation that fall largely on different groups of people – legitimate pain patients bear the costs, while people at risk of becoming illicit drug users experience the benefits. Optimal policy would focus on reducing diversion and theft.³

³Early efforts to rein in the negative effects of the expanded use of prescription opioids focused on this channel – shutting down illicit ‘pill mills’ via DEA busts, preventing pharmacy theft, and reducing the supply of unused pills in medicine cabinets that might be stolen by friends or family members via drug “take back” programs (Smith, 2013; Kolodny et al.,

However, recent concerns have shifted to whether opioids are being appropriately used in the course of medical care. An influential medical literature review published in 2015 by [Chou et al.](#) stated that “accumulating evidence supports the increased risk for serious harms associated with long-term opioid therapy, including overdose, opioid abuse, fractures, myocardial infarction, and markers of sexual dysfunction.” The study also noted that because no randomized controlled studies have been conducted for opioid use that extends past one year, “reliable conclusions about the effectiveness of long-term opioid therapy for chronic pain are not possible.” As such, efforts to reduce opioid-related adverse events have shifted towards more directly targeting the pain management system, and pain patients are seen as most likely to be at-risk. Under this newer understanding, it is possible that doctors and patients are not jointly and optimally trading off the costs and benefits of opioid therapy. On the patient side, time inconsistency and present bias or incomplete information might lead to an irrational use of and the development of a dependence or addiction to these drugs. On the doctor side, agency problems with respect to pharmaceutical company perks, insurer reimbursement systems, or hospital satisfaction surveys might incentivize over-prescription of opioids, even as it may be counter to the long-term welfare of many individual patients ([Thomas et al., 2015](#)).^{4,5} Pharmaceutical manufacturer Purdue Pharma, which makes the opioid pain reliever OxyContin, pled guilty in 2020 to criminal charges including conspiracy to aid and abet doctors in dispensing medication “without a legitimate medical purpose” ([U.S. Department of Justice, 2020](#)).

This paper makes several contributions to this debate. First, I establish that an important early policy lever for addressing societal problems brought about by increased opioid use – the implementation and push towards greater use of state-level Prescription Drug Monitoring Programs – achieved key policy goals of reducing opioid abuse and overdose deaths. However, I also document, using a rich data environment comprising medical claims linked to administrative work records for millions of patients, that there have been unintended consequences with considerable welfare costs from the reduction in opioid availability that has been the result of these efforts. I find that the reduction in opioid availability alters the medical practice of pain management, and patients and their doctors

[2015; SAMHSA, 2012](#)). PDMP regulations also focused on catching illegitimate pain patients, such as “doctor shoppers” who visited multiple clinicians in order to obtain drugs for street sale.

⁴Legislative attention has been focused on both pharmaceutical company perks and hospital patient satisfaction surveys. In 2014 Senators Grassley and Feinstein sent a letter to Centers for Medicare and Medicaid Services complaining that the use of HCAHPS satisfaction surveys in Value Based Purchasing, which have 3 questions out of 20 on pain control, was distorting incentives for hospitals and physicians towards over-prescription of pain medications. In 2012, Senator Grassley requested information from Medicaid which demonstrated that the highest prescribers of certain abusable prescription drugs were also those receiving large payments from the pharmaceutical companies that make them.

⁵Insurance limitations on reimbursements, e.g., for physical therapy and mental health, are well known, and physician responses to recent opioid restrictions have highlighted “access and insurance coverage limitations for non-pharmacologic treatments” as key concerns when implementing opioid reductions. CMS recently made cuts to the reimbursement structure for some interventional pain management procedures due to lack of evidence for their efficacy and concerns about overuse; the American Society of Interventional Pain Physicians responded by stating that “this may end up driving patients to receive more opiates, drive patients to expensive points of care (the hospital) and have patients [undergo] unnecessary spinal surgery, which is much more expensive than the care offered by interventional pain physicians.”

substitute towards more costly, and possibly inferior, forms of chronic pain management; on net, episodes of chronic pain grow more expensive to treat. Additionally, I use a predictive machine learning algorithm to identify a cohort of individuals who are likely to suffer from chronic pain conditions, and find that there is markedly increased missed work under workers' compensation and short-term disability, indicating their pain is more poorly managed after these reforms. This suggests that an increase in pain and a reduction in function has resulted from reduced access to pain management with opioids, and that substitution towards alternate pain management has not been sufficient to fully compensate.

I use these empirical estimates as inputs to a rough welfare calculation, finding that reducing opioid use results in a tradeoff wherein \$8.2–\$10.9 billion per year in increased costs from inpatient and outpatient medical spending plus lost wages are traded off against \$4.8–\$7.6 billion per year in benefits from reduced addiction and lives saved from opioid overdose. However, I find evidence that the groups experiencing benefits from regulation are not the same as groups experiencing harms. I first exploit the linked claims and work data to show that the individuals with chronic pain who experience harms, i.e. loss in work productivity and function, do not appear to be the same cohort of individuals who see benefits through declines in opioid abuse episodes. I also examine patterns of benefit and harm demographically across the medical and mortality datasets, showing that the groups being targeted by, and benefiting from, regulation are broadly but not entirely overlapping demographically. These results are suggestive that social tradeoffs of permissive opioid prescribing are significant, and a move towards more restrictions in prescribing may have exerted an important externality on chronic pain patients.

As discussed in [Garthwaite \(2012\)](#), an increasing enthusiasm for comparative and cost effectiveness research, and in particular a shift towards its use in deciding reimbursement rates and insurance coverage (e.g., “value based care”) has typically neglected to include important broader nonmedical and economic benefits of medical technologies, especially labor supply. This is in part due to a scarcity of evidence on those effects stemming from the fact that it is difficult to identify exogenous variation in the use of medical technology; as such, few studies have directly estimated the economic or labor supply benefits of medical technologies or pharmaceuticals. (Exceptions include [Garthwaite \(2012\)](#), [Powell and Seabury \(2014\)](#), and [Berndt et al. \(2000\)](#).) Pain medications represent a major area of concern: they were the fourth largest by-volume class of drug sold in the United States in 2018, and are dispensed to about 55 million unique patients per year, down from about 75 million at peak levels in 2011–2012 ([Aitken, 2013, 2019](#); [Centers for Disease Control, 2019b](#)).

A few recent studies have looked specifically at the relationship between opioids and labor supply ([Park and Powell, 2020](#); [Savych, Neumark and Lea, 2019](#); [Franco, Wagner and Whaley, 2019](#); [Deiana and Giua, 2018](#)). Identification of the labor supply effects of opioid policies is complicated by the

interlocking nature of economic factors and the opioid crisis broadly, as has been explored in recent work about economic correlates with rising “deaths of despair” (Case and Deaton, 2017; Krueger, 2017). Work focusing on illicit opioid abuse such as Park and Powell (2020) tends to find that illicit opioid availability and higher rates of misuse are associated with reduced labor market supply. The focus of this study is complementary but different. By borrowing several useful supervised and unsupervised machine learning methods and leveraging my very rich data environment to identify and characterize the changing experiences of chronic pain patients under increasingly restrictive prescribing regimes, I am able to shed light on the internal as well as external tradeoffs between economic and medical costs and benefits for this important and widely-used medical technology.

Additionally, I contribute to the medical and public health literature on the health, wellbeing, and life outcomes of pain patients, and whether they are improved or hindered by long-term use of opioids, with an identification strategy that utilizes plausibly exogenous variation in individual opioid supply to study individual outcomes. Because long-term randomized controlled trials of greater than one year have not been conducted on the use of opioids for chronic pain – and almost all studies on opioids have only lasted between 3 and 16 weeks – studies in this literature, detailed below in Section 2, have been observational or used matched-control techniques, and have been hampered by endogenous selection into chronic opioid therapy which is likely to be correlated with pain, functionality, health, work output, and other life outcomes. My results are generally of the opposite sign to what is often found in these studies. For example, studies of workers’ compensation insurance claims find that the utilization of opioid therapy is correlated with *more* missed days of work after an injury.⁶ This concern likely arises from a positive correlation between length of opioid therapy and days missed driven by a common confounder, severity of the injury; by studying a plausibly exogenous reduction in opioid supply, I show that there is actually a *negative* relationship between use of opioids and days missed in workers’ compensation, and that restrictions in opioid prescribing to injured workers *increases* their time out of work.

The paper proceeds as follows. Section 2 provides background on the medical literature on opioids, and documents the changing landscape of opioid use over the past 20 years. Section 3 describes my data, and Section 4 describes the empirical framework. Section 5 presents results on the policy successes of PDMPs, while Section 6 presents results on the unintended consequences of PDMPs. Finally, Section 7 integrates these results into a back-of-the-envelope welfare calculation. Section 8 concludes.

⁶See, for example, NY Times, *Pain Pills Add Cost and Delays to Job Injuries*: “Workers who received high doses of opioid painkillers to treat injuries like back strain stayed out of work three times longer than those with similar injuries who took lower doses, a 2008 study of claims by the California Workers’ Compensation Institute found. When medical care and disability payments are combined, the cost of a workplace injury is nine times higher when a strong narcotic like OxyContin is used than when a narcotic is not used, according to a 2010 analysis by Accident Fund Holdings, an insurer that operates in 18 states.”

2 Background

The first research to suggest that opioids could be used successfully and safely in the management of chronic, non-cancer pain, without generating significant problems with addiction or abuse, was a non-randomized study of 38 patients who were on opioid therapy for as long as 7 years, published by Portenoy et al. in 1986. Several similar follow-up studies, in combination with a prominent article entitled “The Tragedy of Needless Pain” (Melzack, 1990), which stated that “contrary to popular belief, morphine taken solely to control pain is not addictive,” led to a shifting perception among doctors during the 1990s that pain was widely undertreated and that opioids could provide effective relief. A small number of randomized clinical trials on the effectiveness of opioids for chronic non-cancer pain were conducted during this period, indicating opioids could be used for chronic pain, but they answered a narrow question – over a period of 3-16 weeks, are prescription opioids effective at reducing chronic pain and improving function, compared to a placebo drug (Furlan et al., 2006)?

Guidelines and incentives for physicians and hospitals during this period, as well as the overall regulatory landscape, dramatically shifted towards more-liberal pain management across the board. In 1995 the FDA approved Oxycontin for the management of moderate to severe pain, specifically mentioning its use for chronic pain such as back pain. In 1998, The Federation of State Medical Boards Model Pain Policy was rewritten, stating that “the Board recognizes that controlled substances including opioid analgesics may be essential in the treatment of acute pain due to trauma or surgery and chronic pain, whether due to cancer or non-cancer origins” (Federation, 1998). The policy also articulated that “inadequate pain control may result from physicians’ lack of knowledge about pain management or an inadequate understanding of addiction. Fears of investigation or sanction by federal, state, and local regulatory agencies may also result in inappropriate or inadequate treatment of chronic pain patients.” This policy was strengthened further in 2004, stating that “the state medical board will consider inappropriate treatment, including the undertreatment of pain, a departure from an acceptable standard of practice... and will investigate such allegations” (Federation, 2004).⁷ Similarly, for hospitals, revised 2001 standards from the Joint Commission on Accreditation of Healthcare Organizations emphasized “a patient’s right to pain management” (AAACN, 2001), and introduced the “Pain Scale,” a scale of smiling to frowning faces ranging from 0 (“no hurt”) to 10 (“hurts worst”). The Hospital Consumer Assessment of Healthcare Providers and Systems (HCAHPS) patient satisfaction survey, which affects hospital reimbursement rates as of October 2012, includes 3 questions out of 20 on pain control, such as whether “the hospital staff [did] everything they could to help with your pain” (Zusman, 2012).

⁷The policy also explicitly added that “physicians should not fear disciplinary action from the Board for ordering, prescribing, dispensing, or administering controlled substances, including opioid analgesics, for a legitimate medical purpose.”

While in the early 2000s medical consensus continued to on-balance support the notion that opioids were under-prescribed and “the medical use of opioids does not create drug addicts, and restrictions on this medical use hurt patients” (McQuay, 1999), by the mid-2000s two concerns about the rapid expansion in the use of opioids had emerged. First, a concomitant and largely unanticipated sharp rise in opioid abuse, addiction, overdose, and death suggested that the *safety* of widespread opioid use had been overstated (Ballantyne, 2006; Catan and Perez, 2012). Although the initial focus of opioid harm mitigation was on preventing diversion and abuse of opioids by non-medical abusers (e.g., through stealing drugs from others’ medicine cabinets, pharmacy theft, diversion via pill mills, etc.) increasing attention has been paid to abuse, “iatrogenic addiction,” and overdose among medical users (Kolodny et al., 2015; Kaplovitch et al., 2015; Ballantyne, 2012; Chou et al., 2009, 2015). However, attempts to estimate how many cases of opioid use disorder are created via legitimate medical use are hampered by the fact that a clear consensus on the definition of problematic opioid use and abuse among prescribed patients does not exist, and in practice the distinction can be very blurry (Ballantyne and Shin, 2008; Jones, Paulozzi and Mack, 2014).⁸ One literature review (Vowles et al., 2015) found that estimated rates of problematic opioid use across studies are quite broad, ranging from <1% to 81%.⁹

Second, the *effectiveness* of long-term opioid therapy for chronic pain also began to be called into question. Although there is relative consensus that opioids are effective at managing short- and medium-term pain (Furlan et al., 2006; Chou et al., 2009; Noble et al., 2010), due to the lack of randomized clinical trials exceeding 16 weeks, observational epidemiological studies are influential as the best-available evidence on long-term chronic opioid use (usually defined as exceeding 90 days) (Rosenblum et al., 2008; Ballantyne, 2012; Frieden and Houry, 2016). A 2006 study using rich data on Danish adults (Eriksen et al., 2006) studied a group of patients who reported chronic pain, utilizing non-opioid users as a control group for opioid users. Self-reported pain, work capacity, and total medical utilization were compared between the two groups, and opioid users fared worse on all outcomes. Although the study cautioned that “because of the cross-sectional nature causative relationships cannot be ascertained,” it also concluded that “it is remarkable that opioid treatment of long-term/chronic non-cancer pain does not seem to fulfil any of the key outcome opioid treatment goals: pain relief, improved quality of life and improved functional capacity.” A number of other studies assessing opioid efficacy utilize a similar observational methodology, and find that opioid use is associated with increased medical expenditure and more missed work on workers’ compensa-

⁸Nonmedical prescription drug use is most frequently defined as “use without a prescription, or use that occurs simply for the experience or feeling the drug causes.”

⁹In the Marketscan medical claims data studied in this paper, 1.1% of enrollees prescribed oxycodone were later observed with a claim relating to opioid use disorder, as defined below. This figure is 0.64% when considering enrollees prescribed any opioid.

tion and disability (Braden et al., 2012; Manchikanti et al., 2011; Vogt et al., 2005; Mahmud et al., 2000; Morris et al., 2015; Sjogren et al., 2010; Franklin et al., 2008; Gross et al., 2009; Volinn, Fargo and Fine, 2009). However, these studies all suffer from major concerns that selection into opioid therapy is endogenous to the severity and unmanageability of pain, as well as correlated with other characteristics such as poor mental health that might determine outcomes, and hence that the negative association between opioid therapy and health and work outcomes may not be causal.

A 2008 review of the use of opioids for analgesia summarized the clinical situation as follows: “The past several decades in the United States have been characterized by attitudes that have shifted repeatedly in response to clinical and epidemiological observations, and events in the legal and regulatory communities. The interface between the legitimate medical use of opioids to provide analgesia and the phenomena associated with abuse and addiction continues to challenge the clinical community, leading to uncertainty about the appropriate role of these drugs in the treatment of pain” (Rosenblum et al., 2008). The 2015 review (Chou et al., 2015) was more pointed, stating “the lack of scientific evidence on effectiveness and harms of long-term opioid therapy for chronic pain is clear and is in striking contrast to its widespread use for this condition ... Reliable conclusions about the effectiveness of long-term opioid therapy for chronic pain are not possible due to the paucity of research to date.”

Continuing uncertainty is reflected in persistent controversy among doctors, pain practitioners, and public health professionals about the role of opioids in society; some doctors advocate reconsideration about whether opioids should be indicated for most chronic pain patients (McCance-Katz et al., 2012).¹⁰ The March 2016 CDC opioid guidelines recommending major changes in opioid prescribing practices relied heavily on the finding of Chou et al. that no evidence exists to support long-term chronic opioid therapy. In a statement accompanying the guidelines, CDC Director Thomas Frieden reflected shifting attitudes: “what we want to just make sure is that doctors understand that starting a patient on an opiate is a momentous decision. The risks are addiction and death, and the benefits are unproven” (Frieden and Houry, 2016; Bernstein, 2015). Yet the guidelines, and process surrounding the development of the accompanying evidence review, have also provoked controversy and concern among doctors and patient advocacy groups (Anson, 2015; McCarberg, 2016; Madara, 2015; Alford, 2016).

The result of this shifting understanding of the most appropriate role of opioids in society has been a series of regulatory changes that have resulted in a decline since 2011 in total opioids dispensed in the US, as seen in Figure 1. A primary initial policy lever was the introduction of Prescription Drug Monitoring Program laws at the state level, which are detailed below and which are stud-

¹⁰Increasing attention is also being focused on opioid prescribing for acute pain; in 2016 Massachusetts became the first state to limit initial opioid prescriptions, to a maximum of 7 days.

ied in this paper.¹¹ Amidst optimism that these policy levers are reducing opioid abuse, overdose, and death, there have also been increasing reports of difficulty accessing pain medications by chronic pain patients (Gleason et al., 2014; Hoffman, 2016; Centers for Disease Control, 2019a; Madara, 2020).

This debate over the appropriate societal role of opioid pain relievers in the treatment of pain remains unresolved, and is the subject of discussion and debate by clinicians, researchers, and policymakers. A landmark randomized clinical trial of the use of opioids versus non-opioid medications to treat 240 veterans with chronic back pain or osteoarthritis pain, which notably extended to 12 months, found that the use of opioids did not result in significantly improved pain-related function compared to non-opioid medications (Krebs et al., 2018). Busse et al. (2018) conduct a meta-analysis intended to update previous meta-analyses such as Chou et al. (2015), reviewing 96 randomized clinical trials of opioids for chronic noncancer pain; they find the use of opioids is associated with a statistically significant reduction in pain and increase in physical functioning, but that the magnitude of the effect is small. Finally, Nadeau, Wu and Lawhern (2021) conduct an analytic review of the clinical literature, finding substantial but not definitive evidence that opioids are effective in treating chronic pain; they also highlight the continuing absence of comparative effectiveness studies that prove effectiveness of non-pharmacologic substitutes. Reflecting these ongoing concerns, the CDC recently released a commentary clarifying its 2016 Guideline (Dowell, Haegerich and Chou, 2019), encouraging some relaxation of overly-strict interpretations; in 2020 the American Medical Association released a new statement “welcom[ing] the [2019] clarification,” but calling for the CDC to further revise its approach, stating that “it is clear that the CDC Guideline has harmed many patients.... patients suffering from pain increasingly view themselves as collateral damage. A dedicated effort must be made to undo the damage from the misapplication of the CDC Guideline.”

3 Data

3.1 Drug Enforcement Administration Automated Reports and Consolidated Orders System

Data for this study comes from four main sources. First, I digitize reports from the Drug Enforcement Administration’s Automated Reports and Consolidated Orders System (ARCOS).¹² The ARCOS system tracks the flow of DEA controlled substances from their point of manufacture, through

¹¹Several other regulatory changes have contributed. In 2010, Purdue Pharma switched to a new abuse-deterrent formulation of Oxycontin, intended to provide equivalent pain relief to the old formulation, but prevent street abuse by being resistant to crushing and snorting. In 2011, the DEA engaged in a series of crackdowns on pill mills, especially in Florida. In 2013, the FDA blocked the introduction of non-abuse-deterrent generic versions of Oxycontin from entering the market.

¹²These reports are publicly accessible on the DEA website at https://www.deadiversion.usdoj.gov/arcos/retail_drug_summary/index.html; I have published the data and a package to process the yearly PDFs on <https://github.com/akilby/arcos>.

commercial distribution channels, to the point of sale/distribution (by hospitals, retail pharmacies, practitioners, mid-level practitioners, and teaching institutions). The ARCOS reports contain data on all Schedule II and select Schedule III substance sales at the quarter-state level; in particular, these reports cover the sale of all major opioid pain relievers broken out by active ingredient (codeine, fentanyl, hydrocodone, hydromorphone, meperidine, methadone, morphine, and oxycodone).¹³ An advantage of this dataset relative to all other data sources on opioid prescribing is that it contains a census of all U.S. prescribing and thus represents a complete accounting of all opioids available in the U.S. net of any substitution between sources.¹⁴

3.2 Marketscan Commercial Claims and Encounters / Health and Productivity Management

Second, I utilize the Truven (now IBM) Marketscan Commercial Claims and Encounters (CCAЕ) plus linked Health and Productivity Management (HPM) database from 2005 to 2012. This individual-level panel dataset is constructed from health insurance claims obtained from large employers, and has grown from 1,291,067 enrollees in 2005 to 3,964,364 enrollees in 2012.¹⁵ Broadly, the individuals in a given year of the Marketscan data can be considered to represent the population of working Americans who have employer-sponsored insurance, or about 59% of the total population. The CCAЕ dataset contains rich clinical indicators: detailed prescription claims, clinical utilization, and expenditures, including diagnoses, procedures, type of provider, etc. Additionally, the CCAЕ data is linkable via unique enrollee ID to the HPM databases, which contain administrative workplace data on absences, short-term disability, long-term disability, and workers' compensation data. The median enrollee appears in the dataset for about three years.

Because the Marketscan sample is obtained from large employers, individuals usually move into and out of the sample in groups based on whether their employer is providing data to Marketscan for that year. This generates concerns that non-random selection of individuals into and out of the

¹³Schedule II and III drugs are both defined as having "a currently accepted medical use in treatment in the United States," and a "potential for abuse and dependence." Fentanyl, hydromorphone, meperidine, methadone, morphine, and oxycodone are Schedule II, while codeine and hydrocodone are Schedule III. Schedule III drugs are distinguished from Schedule II as less abusable, and less likely to produce physical and psychological dependence. In practice, the distinction is important because Schedule III drugs are easier to prescribe and obtain: they can be refilled and the prescription can be called in, while Schedule II drugs cannot. They also have been less of a focus on enforcement by the DEA, leading to greater practitioner comfort with their use.

¹⁴However, the data is also very noisy, and this is likely a significant source of measurement error in my ARCOS analysis. First, the ARCOS system has some internal inconsistencies. For example, Report 2 contains quarterly drug distribution in each state in *grams*, and Report 3 contains quarterly drug distribution in each state in *grams per 100k population*. Despite the fact that each quarter-state observation should only differ by a constant (100k population as a divisor), these numbers in Reports 2 and 3 – which were generated for me simultaneously by the DEA – do not always agree. The ARCOS reports are also regularly updated and occasionally archived historical reports are re-run (indicated by a new Run Date). This often changes the dataset slightly as well.

¹⁵Overall, 7,131,438 total individuals appear in the database.

sample could bias my estimates. In order to handle this concern, my baseline empirical specification will utilize individual-level fixed effects, thus estimating only within-person variation. Results using individual fixed effects are very similar to results using entry-cohort fixed effects. I explore robustness concerns relating to the sample and choice of fixed effects in [Appendix D](#).

For most analyses, I extract an outcome variable from the CCAE claims or HPM data and aggregate it to the individual-quarter level. For example, a main outcome variable, “quantity of oxycodone consumed by an individual in a quarter,” is derived from the prescription drug flat file, which contains a list of every prescription filled, identified by National Drug Code (NDC) number (including date filled, intended days supply, and quantity). Truven produced a datafile called RED BOOK to map NDC numbers onto active ingredients, number of dosage units, and strength per dosage unit. To identify oxycodone prescriptions, I first process the RED BOOK datafile to identify all NDC numbers with oxycodone as the active ingredient. I then merge RED BOOK onto the CCAE prescription drug flatfile to identify oxycodone prescriptions. Next, I standardize the unit of measurement, strength per unit, and quantity measures in order to calculate the milligrams of active ingredient in each prescription. (For combination prescriptions with multiple active ingredients, only milligrams of the opioid in the prescription were included.) Finally, I total the active milligrams of oxycodone in all oxycodone scripts for each individual in each quarter. A similar procedure is utilized for outcome variables derived from inpatient and outpatient claims, as well as workers’ compensation and disability data.

As can be seen in [Table 2](#), in a given year approximately 25% of enrollees in the Marketscan sample will fill a prescription for an opioid.¹⁶ This is in conformance with national data for this period, during which about 75 million patients receive narcotics prescriptions in a given year ([Aitken, 2013](#)). In this study I focus primarily on the prescribing of Schedule II opioids, because they are used for moderate to severe pain and because the reduction in their use has been a primary target of opioid crackdowns. Oxycodone, a Schedule II drug, is the highest-sold opioid by volume, and is also the drug most frequently targeted by these crackdowns, due to the fact that it has typically been considered the most abused and abuseable commonly-prescribed opioid. The highest-volume Schedule III opioid is hydrocodone; it is the opioid patients are most commonly exposed to in a given year, as can be seen in [Table 2](#). This difference in frequency versus total volume is because hydrocodone is more typically used for acute, less-severe, incidental, and shorter-term pain, and thus is prescribed in low and short doses, whereas oxycodone is typically used for chronic or more-severe pain, and thus is prescribed in higher doses and for longer-term use.¹⁷

¹⁶Note that these sample statistics are reported only for the sample of enrollees in the 38 states that are in my study sample, which is described in detail in [Section 4.1](#) below.

¹⁷Although oxycodone is a slightly stronger opioid than hydrocodone, with a MME of 1.5 compared to a MME for hydrocodone of 1, this difference in utilization patterns is mostly due to regulatory factors. Hydrocodone formulations were,

Opioids naturally generate tolerance in their users, and most long-term opioid patients will experience escalating doses, such that a typical dose for an opiate-naive user will be lower than a typical dose for an experienced user (and in fact a dose for an opiate-experienced user may be sufficiently high that it would be fatal to an opiate-naive user). During the period of this study (until 2012) there was not a medical consensus on whether doses should be allowed to escalate to very high levels, or rather should be capped at some maximum; guidelines have been increasingly forceful about establishing a maximum for most patients.¹⁸ In the MarketScan prescription claims dataset, I group prescriptions into therapy episodes, where multiple prescription claims are considered to be part of the same episode of opioid therapy if they are for drugs with the same generic name and commence within 3 times the number of days supply of the last prescription filled with that same generic name. Most prescriptions for oxycodone and hydrocodone – over 80% – are initial prescriptions with no follow-on prescriptions, and last less than 15 days. Only a small fraction of opioid therapy episodes (4.5% for oxycodone, 3.3% for hydrocodone) last longer than 90 days. As shown in Table 2, the average Morphine Milligram Equivalent (MME) per day prescribed of oxycodone escalates over time as the duration of the therapy episode increases; the same increase is not observed for hydrocodone because in all formulations available during this period the maximum safe dose was constrained by the maximum safe dose of other active ingredients, usually acetaminophen.

The end result of these prescribing patterns is an extremely right-skewed distribution, where most observations in an enrollee-quarter for oxycodone or hydrocodone are zero, but the tail of the distribution is also very long. My baseline specification will utilize a concave transformation of this outcome variable (specifically, $\ln(\text{MME in quarter} + 1)$), which preserves zeros but reduces the influence of outliers. As can be seen in Tables 1 and 6, the average MME per enrollee per quarter of all Schedule II opioids in the MarketScan sample (60.99 mg) is lower than the MME per resident per

until recently, only approved in low doses and in combination with other drugs such as acetaminophen, which, due to its liver toxicity, is considered to provide some abuse deterrence. The presence of other ingredients such as acetaminophen means that hydrocodone products are less safe for long-term use, and cannot be escalated to higher dosages. All hydrocodone formulations (e.g., Vicodin, Lortab) were Schedule III throughout the study period, whereas all oxycodone formulations (as well as morphine, fentanyl, hydromorphone, etc.) were Schedule II, and due to being in Schedule III, hydrocodone was easier to prescribe and obtain than oxycodone. Hydrocodone was rescheduled to Schedule II in August 2014, in large part due to the controversial 2014 FDA approval of Zohydro, a hydrocodone formulation which comes in much higher doses, does not contain any additional active ingredients such as acetaminophen, and which does not have any abuse deterrent technology.

¹⁸Older pain management guides state that “because of inpatient variability in physical dependence, opioid tolerance, subjectivity of pain, and biopsychosocial influences, exact opioid dosing with consistently accurate equivalency tables (for conversion among opioids) are lacking. Accordingly, opioids have no well-defined maximum dosage to achieve appropriate therapeutic benefit. Dosage escalation is commonly used until adequate pain relief is achieved or adverse effects preclude further escalation. One panel has stated that a ‘reasonable definition for high dose opioid therapy is >200 mg daily of oral morphine [sulfate] (or equivalent), based on maximum opioid doses studied in randomized trials.’ (Nerenberg and Fudin, 2010) In 2012, Washington State set a threshold of 120 MME per day, above which primary care physicians are required by law to consult with a pain management specialist. In Massachusetts, 2015 guidelines recommend that any patient receiving over 100 MME per day begin tapering to a lower dose. Finally, the 2016 CDC guidelines recommended a maximum dose of 90 MME per day, with increased prescriber caution above 50 MME per day, as discussed previously.

quarter in the ARCOS data (106.7 mg), with the lower figure likely reflecting the fact that the Marketscan sample does not represent the Medicare-covered elderly or the Medicaid-covered poor, who have greater utilization of these drugs, and also would not capture any illegal or suspicious high-volume diversion, e.g., the operation of pill mills, which do not usually accept insurance (Ailes et al., 2015; Hackbarth, Christianson and Miller, 2015).

Other key outcome variables from this data source include inpatient and outpatient spending, number of days absent from work under workers compensation and short term disability programs, and whether the patient is identified as having an episode of opioid use disorder. I also aggregate all spending identifiable in the Marketscan data – inpatient, outpatient, and drug spending, as well as medical costs and lost wages under workers compensation, and lost wages under short term disability, into a total costs measure.

3.2.1 Identification of chronic pain patients and distinct episodes of care

I use this detailed medical and administrative work productivity data to tell a rich, data-driven story about the experiences of pain patients under increasing opioid prescribing restrictions. Specifically, I identify individuals who are likely to need pain management due to chronic pain, and document changes in their inpatient and outpatient care, as well as their likelihood of taking absences from work.

Identifying chronic pain patients is non-trivial; no single diagnosis or procedure code identifies a patient suffering from chronic pain conditions, which range from back pain and headache to fibromyalgia and arthritis. Rather, a constellation of diagnosis and procedure codes characterizes clusters of experiences that are usually associated with pain.

In order to identify a cohort of patients most likely to experience tradeoffs related to reductions in opioid prescribing, I utilize a machine learning algorithm, a tree-based gradient boosted model (Friedman, Hastie and Tibshirani, 2009), to identify correlations between chronic opioid therapy and a rich set of clinical covariates. I use this fitted model to generate a predicted subsample of individuals most likely to suffer from chronic pain conditions. I describe the approach to model training and construction of this chronic pain cohort in Appendix E. I present summary statistics of individuals in this chronic pain subsample, compared to the full sample, in Table E2. For many analyses, I focus on the experiences of this smaller cohort, as it would not be expected that outcomes like diminished work productivity would accrue to patients only receiving opioids for acute pain.

I further characterize the experience of these chronic pain patients by segmenting their medical history into distinct episodes of care, and characterizing these episodes of care using an unsupervised natural language processing topic modeling algorithm, latent Dirichlet allocation, which classifies data into topics where each observation in a collection is substantially similar. This ap-

proach allows the identification of interpretable episodes of care relating directly to the treatment of pain conditions (e.g., back surgery); episodes of care that may have pain as a side effect (e.g., cancer); and unrelated episodes of care that occur at other points during the patients' time in the sample. This approach is described in [Appendix F](#).

3.3 National Vital Statistics System Multiple Causes of Death Microdata

Finally, I utilize CDC National Vital Statistics System (NVSS) Multiple Cause of Death Microdata. This dataset contains a record of every death in the United States from 2001 to 2013, including underlying cause of death and up to 20 additional causes of death listed on the death certificate, demographic data such as race, gender, age, and education, and geographic identifiers at the county level. I follow the public health literature in identifying drug overdoses caused by prescription opioids and heroin using both the underlying and multiple cause of death fields.¹⁹ Heroin-related drug overdose deaths are identified if the multiple causes of death fields contain code T40.1, and opioid-related drug overdose deaths are identified if they contain codes T40.2 (morphine, oxycodone, hydrocodone), T40.3 (methadone), T40.4 (synthetic narcotics – Fentanyl, Propoxyphene, Meperidine, Buprenorphine), or T40.6 (“other and unspecified narcotics,” often used when only “opioid overdose” is indicated by a medical examiner ([Slavova et al., 2015](#))). For my prescription overdose death results, I will in most cases present estimates for T40.2 alone, because it is comprised of the Schedule II drugs for which I find a first-stage prescribing effect. For welfare calculations I will combine all opioids and heroin, i.e. T40.1-T40.4 and T40.6, in order to estimate the aggregate impact on opiate and opioid deaths, which will include any offsetting dynamics as users substitute between different types of opioids/opiates.²⁰

4 Empirical Framework

I exploit state-level variation in the timing of the introduction of Prescription Drug Monitoring Program (PDMP) laws to study the impact of PDMP introduction on opioid prescribing, as well as establish a causal linkage between a reduction in opioid prescribing and a range of other health and labor outcomes of interest.

¹⁹I identify drug overdoses using Underlying Cause of Death codes X40-X44, X60-X64, X85, and Y10-Y14.

²⁰Multiple Cause of Death categories are not mutually exclusive – it is possible to have both heroin and prescription opioids contribute to a death, and for both codes to be listed in the 20 multiple cause of death fields. This overlap has increased considerably over the past decade, mostly after the period of study of this paper, as can be seen in [Figure 2](#). Over years in my sample, about 12% of opioid overdoses also had heroin listed on the death certificate.

4.1 Institutional details of Prescription Drug Monitoring Programs

As discussed in Section 3.1, the Drug Enforcement Administration (DEA) monitors the flow of controlled substances in the United States, but active monitoring and enforcement largely ends at the second-to-last step in the distribution chain: the sale of those substances from manufacturers/distributors to pharmacies, hospitals, and practitioners. Monitoring the last step of the chain, from pharmacies to patients at the retail level, is largely left to the states. State-level programs to monitor the flow of these substances date back as early as 1973, and a common approach during this early period used special triplicate prescription pads to report the prescribing of every controlled substance. Between 1973 and 2003, 12 states implemented some kind of controlled substances monitoring; these efforts focused on providing information to law enforcement or licensing boards for the purposes of detecting illicit behavior among patients and doctors.

As opioid prescribing began to grow rapidly in the early 2000s after physician and hospital standards on pain management were liberalized, concerns about greater oversight spurred increasing attention to state-level monitoring. In 2002, the DOJ started a grant program to states to support the implementation of state Prescription Drug Monitoring Programs through the Harold Rogers Prescription Drug Monitoring Program (HRPDMP), and began cooperation with the National Alliance for Model State Drug Laws (NAMSDL), which published the first Model Prescription Drug Monitoring Program Act for adoption by state legislatures in 2003.

The HRPDMP maintained a concern with law enforcement, but during this period states also began to recognize that PDMPs could also serve to improve pain management practice and the coordination of care by providing information to physicians and pharmacists about their patients. Nevada introduced the first PDMP that was partly intended to improve the practice of pain management via providing access to the database to physicians and prescribers in 1997. The 2003 NAMSDL Model PDMP Act included language that the PDMP database should be accessible to prescribers. In 2005, this shift in focus towards using PDMPs in the practice of medicine (rather than primarily facilitating law enforcement) was codified in federal guidelines passed as a part of the National All Schedules Prescription Electronic Reporting Act (NASPER), which also included grant funding for the implementation of state PDMPs. Due to the NAMSDL model law, most states adopting after 2005 adopted similar legislation, including an electronic database and online access for physicians and pharmacists at the point of care, and the focus of those PDMPs was on altering the practice of medicine.^{21,22}

As such, states implementing PDMPs during this period were implementing relatively standardized laws, and the introduction of a PDMP can be thought of as a natural experiment. In this natural

²¹These later-adopting states even often utilize the same software, from one of a small handful of vendors.

²²In some recently-adopting states the relationship between law enforcement and the PDMP is actively non-cooperative, and access requires probable cause or a subpoena; see e.g., *Oregon PDMP v. United States Drug Enforcement Administration*.

experiment, physicians simultaneously gain access to information about their patients, and also begin undergoing monitoring themselves. PDMPs can thus be thought of as a bundled intervention which might affect opioid prescribing and distribution in several ways. First, PDMPs provide information on patients which can be used to improve care, such as if a patient is being prescribed multiple medications by multiple doctors which are in combination contraindicated due to a drug interaction. Second, by increasing oversight of the last point in the distribution chain, they should reduce obviously illegal diversion by pharmacists, physicians (pill mills), and patients (doctor shoppers). (Patients might be caught by law enforcement if they are engaging in illegal behavior, but there are other behaviors, such as a violation of a patient contract specifying the patient may only visit one doctor, which the PDMP can reveal.) Finally, in conjunction with shifting norms and standards towards the use of opioids, the fact of being monitored by PDMPs should induce doctors to move their own prescribing practices more in alignment with “best practice,” and in general err on the side of prescribing fewer rather than more opioids.²³ In practice, exactly how PDMPs affect prescribing practices is an empirical question that will be explored in Section 6.

I exclude the 12 states that implemented some kind of controlled substances monitoring program prior to 2003 from my analysis (see Table A2). Institutionally, the highly variable nature of early controlled substances monitoring approaches at the state level suggests that states with early PDMPs are likely to be qualitatively different and are thus poor controls for the study of opioid prescribing during the period of study. In particular, assigning a “date of implementation of PDMP” to these states is difficult; early controlled substances monitoring programs usually did not meet all the criteria of a PDMP that this study considers (especially physician access), and states with early programs tended to phase in the characteristics associated with a PDMP that states in the 38-state sample usually adopted all at once, while simultaneously experimenting with more aggressive policies like proactively reaching out to doctors about their patients.²⁴

Figure A2 demonstrates that these institutional differences indeed translated into notably different prescribing patterns in early-adopting and late-adopting states, both in level and in trend. Early controlled substances monitoring states saw significantly fewer opioids prescribed and experienced a slower growth in opioid prescribing during the study period, and appear to have partially substituted for lower Schedule II opioid usage with greater hydrocodone usage, which is higher in level

²³Alpert, Dykstra and Jacobson (2020) document that “must access” PDMPs, which require physicians to access the PDMP prior to writing an opioid prescription, operate both through an information channel and via imposing “hassle costs” that increase the cost of prescribing opioids to all patients by consuming valuable clinician time for every prescription written. Hassle costs imposed by states may in some sense formalize the softer push towards new norms of earlier programs.

²⁴For example, NAMSDL lists Illinois as having implemented its controlled substances program variously in 1961 or 1999; the state lists the date of implementation as 1984. And in a 2008 press release announcing reforms that brought its controlled substances program into conformance with national guidelines and thus finally meeting the definition of a PDMP, Illinois also called itself a leader in the aggressive monitoring of controlled substances.

and in growth rate in early adopting states. Other literature on states which implemented early versions of prescription monitoring have similar findings: [Reisman et al. \(2009\)](#) found lower overall opioid prescribing in the 1997-2003 period in the states that had already adopted PDMPs at that point; [Alpert et al. \(2019\)](#) found that triplicate pad states saw significantly reduced distribution of OxyContin in particular, and subsequently had lower overdose death rates. Starting in 2012/2013 many of the states in the 38-sample began to strengthen their PDMPs, in particular by mandating that physicians access the PDMP prior to dispensing any controlled substance, a reform that has been advocated as a “best practice” for PDMPs, so focusing on the period after 2002 to 2013 provides a relatively clean experiment.

Public health and economic studies evaluating the effectiveness of PDMPs thus far have been mixed: most studies find support that PDMPs alter prescribing patterns, but evidence on reductions in rates of opioid misuse, abuse, and overdose deaths is weaker. Later studies considering PDMPs with strengthened provisions, such as requirements that clinicians access the database prior prescribing, are more likely to find effects ([Maclean et al., 2020](#); [Meinhofer, 2018](#)). [Paulozzi, Kilbourne and Desai \(2011\)](#) found that between 1999 and 2005, for states with operational PDMPs, Schedule II prescribing declined, hydrocodone prescribing increased, and overdoses were not affected. [Twillman \(2006\)](#) documented similar patterns for Schedule II and Schedule III prescribing, and also found no impact on measures of abuse, overdose, and death. [Reifler et al. \(2012\)](#) found that between 2003 and 2009 states with PDMPs had a lower rate of increase in indicators of opioid abuse and misuse (but did not consider deaths). [Rutkow et al. \(2015\)](#) found that the introduction of Florida’s PDMP in 2011 was associated with significant declines in prescribing compared to a neighboring state, Georgia, and [Delcher et al. \(2015\)](#) found this reduction in prescribing corresponded with a reduction in overdose deaths. More recently, [Moyo et al. \(2017\)](#) find that PDMPs reduce volume of opioids dispensed in the Medicare program, [Bao et al. \(2016\)](#) find similar effects in a survey of prescribing provided in ambulatory care settings, and [Buchmueller and Carey \(2018\)](#) and [Bao et al. \(2018\)](#) find that “must access” PDMPs reduce inappropriate patterns of prescribing and indicators of patient opioid misuse in the Medicare program and a commercial claim dataset, respectively. [Patrick et al. \(2016\)](#) and [Pardo \(2017\)](#) find that PDMPs reduce overdose deaths, especially robust PDMPs (with more intense regulatory requirements that employ PDMP “best practices”). Overall, PDMPs appear to affect prescribing patterns across a variety of settings, but evidence on abuse and overdose outcomes is still somewhat inconclusive: [Fink et al. \(2018\)](#) conducted a systematic review of this literature and found inconsistent and weak evidence on the impact of PDMP implementation on outcomes of nonfatal and fatal overdoses.

Following most of the literature above, my primary source for dates of PDMP implementation is

the National Alliance of Model State Drug Laws.²⁵ For most specifications examining the impact of PDMPs on the practice of medicine, I code the initiation of a PDMP according to date of physician access (Table A1). However, there is usually a phase-in period in which pharmacists begin reporting filled prescriptions to the database before physicians gain access, and it is likely that pharmacists and others earlier in the supply chain became aware of, and felt monitored by, the PDMP's implementation before physicians do. Therefore, for the regressions considering surveillance data on the census of all U.S. opioid prescribing (the ARCOS data), I utilize the date PDMP begins collection from pharmacists, which is typically 3-12 months prior to date of access.^{26,27}

The identifying assumption of my empirical strategy is parallel trends in adopting (treatment) and non-adopting (control) states before the implementation of the law; for all my results I will rely heavily on event-study style specifications, described below in Section 5, which estimate leads and lags of the law's implementation, and thus allow for visual verification that the parallel trends assumption holds. I also conduct further informative robustness checks using recent literature on two-way fixed effects estimators with heterogeneous treatment effects in Appendix C. I interpret the introduction of a PDMP as a mechanism that on average reduces individual opioid consumption, and interpret my results on e.g., work absences as causally linked to a reduction in opioid availability. This interpretation requires that there is no other channel by which the introduction of a PDMP might affect work and functionality other than through the availability of opioids.²⁸

²⁵An archived version reflecting the period of study can be found here:

<https://web.archive.org/web/20180419233955/https://namsdl.org/library/580225E9-E469-AFA9-50E7579C1D738E71>; in a few cases internet searches of news articles regarding a PDMP becoming operation yielded slightly different dates, so there are a few discrepancies from this archived version.

²⁶Date of collection is not available for every state; in states where collection date is missing or there is a large discrepancy of over one year between date of access and date of collection, I code the date of collection as the quarter prior to date of physician access.

²⁷Horwitz et al. (2020) et al consider the somewhat inconclusive literature on the effects of PDMPs (detailed below) and show that there is meaningful inconsistency in coding of various states' dates of PDMP implementation. I consider robustness of my results to their legal coding vis-a-vis the more-commonly used NAMS DL database in Appendix A. I also conduct robustness checks to the choice of 2003 as a cutoff, shown in Table B1. Additionally, in two states (OH and NM), several phases of implementation did occur during my sample period (Kasich, 2015; Board of Pharmacy, 2012). In both cases I code using the latter date of implementation, but results are robust to alternate coding using the first date, or to dropping those states from the sample. Because my focus in this paper is on characterizing the impact of a reduction in opioid availability on prescribing and the practice of pain management with the PDMP as a lever that changes opioid availability, rather than on studying the impact of PDMPs per se, I prefer to use the later date of implementation for those two states, as this is when the strengthened version of their PDMP came into effect.

²⁸A slight nuance when considering this identifying assumption is that because PDMPs are a bundled intervention, we might expect that the channels by which opioid supply are reduced are several. In particular, a patient may find his or her opioid supply reduced by their doctor because the doctor found something concerning in their record. They may never be prescribed an opioid because the doctor fears prescribing too many and being flagged in the PDMP. Alternatively, even if they are not a pain patient, if they are an illicit drug user, they may find their opioid use reduced because the overall supply of illicit opioids has been reduced post-PDMP. For these outcomes, I cannot directly detangle whether an increase in absences is due to a reduction in *legitimate* versus *illegitimate* opioid use. I will consider some distributional results which shed light onto the channels in operation, but I am for the most part not able to observe this directly.

5 Empirical Analysis: Policy Successes of PDMPs

5.1 Opioid distribution decreases

I first consider the impact of the introduction of a PDMP on the amount of opioids distributed, utilizing the DEA ARCOS monitoring data for the sale of prescription opioids.

I estimate the following regression models:

$$Y_{st} = \alpha + \gamma_s + \lambda_t + \sum_{\tau} \sigma_{\tau} D_{\tau,st} + \beta X_{st} + \epsilon_{st} \quad (1)$$

$$Y_{st} = \alpha + \gamma_s + \lambda_t + \beta_1 1(PDMP_{st}) + \beta_2 X_{st} + \epsilon_{st} \quad (2)$$

where s is state and t is time (quarter), and $D_{\tau,st}$ are dummy variables for each year before and after the policy is introduced. τ is normalized to 0 in the year before PDMP data collection begins. Y_{st} is milligrams of morphine equivalent sold per state resident. γ_s and λ_t are state and time fixed effects. X_{st} are state controls including unemployment rate and population over 60, and population. Robust standard errors, clustered at the state level, are reported.

Specification 1 utilizes an event-study approach that allows visualization of the impact of the law over time through the estimated coefficients on leads and lags, σ_{τ} . This specification, versions of which I will utilize heavily throughout this paper, allows visualization of the law's impact over time, and also facilitates verification of the parallel trends identifying assumption, by examining coefficients on dummies for the period prior to the law's implementation. Specification 2 is a less-flexible specification which summarizes the mean effect of the introduction of a PDMP using β_1 , the estimated coefficient on a dummy indicator for presence of the law, $1(PDMP_{st})$.

The event study for impact of the introduction of a PDMP on Schedule II distribution, estimated according to Equation 1, is depicted in Figure 3. For Schedule II opioids, there is a decline in dispensing that lasts for about $\tau = 4$ years, before trending somewhat back towards zero. There is no evidence of clear differential pretrends that would indicate policy endogeneity or threaten the empirical framework's internal validity.²⁹ The corresponding point estimate, estimated according to

²⁹If I utilize date of physician access (as is used for the rest of the paper) rather than date of the initiation of data collection from pharmacies, the event study specification shows a slight anticipatory decline in opioid prescription in the year prior to the law's implementation. There are several reasons why prescribing might start to decline earlier in the ARCOS data, and why the same effect is not apparent in the Marketscan prescribing data, as will be seen below. Although PDMPs are not usually high-profile policy initiatives with long lead-times before enactment, they do often have a phase-in period in which pharmacists begin reporting filled prescriptions to the database before physicians gain access, and it is likely that pharmacists became aware of, and felt monitored by, the PDMP's implementation before physicians. Assuming that some diversion occurs at the pharmacy level, this could explain the anticipatory decline in the year prior that shows up only in the ARCOS surveillance data. A related but separate story would be that physicians operating pill mills – who would show up in ARCOS data but not Marketscan because pill mills usually operate cash-only – pay more attention to the passage of PDMP laws, and thus adjust their behavior in advance of the date of physician access, as soon as they are aware of being monitored. These results highlight, as discussed above, that PDMPs are a bundled intervention that reduce prescribing

Equation 2 and displayed in Table 3, indicates that the quantity of opioids prescribed per resident per quarter in a state declines by about 11.15 MME, compared to an overall sample mean of 106.7 MME, representing an overall reduction in opioid dispensing per capita of about 10.4%. Columns (2)–(3) of Table 3 shows that while a substantial fraction of Schedule II prescribing is for drugs other than oxycodone, the entirety of the main effect is explained by a reduction in oxycodone, which is, as discussed above in Section 3.2, the primary target of these laws.³⁰ Finally, Column (4) shows that there is possibly a partial substitution towards hydrocodone prescribing, which is a result that echoes similar results in the PDMP literature review above. Hydrocodone is a weaker opioid that was in this period in a lower DEA schedule (Schedule III) and is typically prescribed in a limited dose; this partial substitution only offsets the overall decline in prescribing by about one-fifth.

Finally, I subject the Schedule II opioid distribution results to a battery of robustness checks, depicted in Table B1. I show results are robust to alternate codings of the legal dates of implementation, and alternate cut-off dates for what is excluded as an ‘early PDMP’ state; this produces different state samples. I show results are robust to inclusion of state-specific linear trends, suggesting differential state pre-trends are not driving the results. Results in columns (8) and (9) for different cut-off dates for being designated an ‘early PDMP’ suggest that there may be heterogeneous treatment effects of PDMP laws, with later implementations more effective. Recent work, discussed in Appendix C, has demonstrated that in many difference-in-differences applications there is significant bias in the two-way fixed effects estimator when there is heterogeneity in the treatment effect and treatment timing varies. This motivates the robustness check in column (11), where I implement the [de Chaisemartin and D’Haultfoeuille \(2020\)](#) estimation approach.^{31,32}

The result that total Schedule II dispensing declines post-PDMP indicates that these laws have had a clear policy success. I will explore in more detail the ways in which Schedule II opioid prescribing and dispensing changed below, in Section 6.1, where I can exploit rich prescription claims to

both by increasing monitoring and scrutiny of physicians and pharmacists, and by providing physicians more information on their patients.

³⁰The relative size of the non-oxycodone Schedule II drugs category, with a sample mean of 47.9 MME as shown in Table 3, is possibly somewhat overstated. This category is dominated by Fentanyl, which is a very potent opioid used for the treatment of cancer in highly opioid tolerant individuals. Because its MME is 75, conversion to morphine equivalents makes fentanyl look higher-volume than it is in practice.

³¹ The only robustness check in Table B1 yielding a substantively different result is the population-weighted regression. This is a result of the fact that Florida, a large state, had an extremely large decline in opioid prescribing coinciding with the introduction of its state PDMP. This was likely due in part to reasons not fully attributable to the PDMP itself: the DEA and state law enforcement undertook a concerted effort during the same time period to raid and shut down pill mills that had been operating prolifically, notoriously supplying much of the southeast’s supply of illicit pills ([Meinhofer, 2015](#)). Florida is an outlier in regards to the level of prescribing attained prior to crackdown as well as the law enforcement scrutiny it received coinciding with PDMP introduction, which made national headlines; this kind of contamination of the PDMP ‘experiment’ by other law enforcement efforts is not present in other states. Dropping Florida from the sample reduces the point estimate to a similar level and significance as the unweighted specification.

³²Results also pass a placebo exercise, not shown, where each state is randomly assigned a date of implementation to my 38-state sample, then estimating β_1 for the placebo PDMP laws 1000 times. My baseline point estimate is lower than 99.2% of estimated β_1 s.

better understand the nature of this decline.

5.2 Opioid overdose deaths decline

I next consider whether deaths indeed declined after the introduction of a PDMP. I estimate the following specifications on the CDC mortality data:

$$Y_{cst} = \alpha + \gamma_c + \lambda_t + \sum_{\tau} \sigma_{\tau} D_{\tau,st} + \beta X_{cst} + \epsilon_{cst} \quad (3)$$

$$Y_{cst} = \alpha + \gamma_c + \lambda_t + \beta_1 1(PDMP_{st}) + \beta X_{cst} + \epsilon_{cst} \quad (4)$$

where Y_{cst} is county-level opioid overdose deaths (T40.2), specified as $\ln(\text{overdose deaths}_{cst} + 1)$ in the baseline OLS specification to account for zeros, and observations are weighted by county population.³³

I document that the reduction in opioid prescribing brought about by the introduction of a PDMP has a substantial and lasting effect on opioid-related overdose deaths. As is shown in Figure 4 and Table B3, Column (1), the introduction of a PDMP resulted in an approximate decline in overdose deaths of around 8% after 1-2 years, and reached 16%-17% in later years.

I conduct robustness checks on this result in Tables B2 and B3. Table B2 conducts a battery of checks related to the choice of PDMP legal coding, similar to those conducted on the ARCOS data. Table B3 considers robustness to different choices of models. In Panel C, I show that the dynamic treatment effect specification is robust to a wide variety of model choices, including a Poisson count model and a linear model estimated on the overdose death rate. However, the static two-way fixed effects estimator is fragile to specification choice. As discussed above and in Appendix C, this bias frequently arises in difference-in-differences applications with variation in treatment timing and heterogeneity in treatment effects.³⁴ This motivates the estimation and reporting of a dynamic average treatment effect for the treated states (ATE/ATT), estimated according to the procedure in [de Chaisemartin and D'Haultfoeuille \(2020\)](#). The procedure produces much more stable estimates of the impact of the law, as can be seen in Panel B in Tables B2 and B3. This yields a summary point estimate of 10.8%, or the impact of the introduction of a PDMP on overdose deaths was a 10.8% reduction in fatalities. This dynamic ATT is also reported in Table 4.

³³I consider robustness to this transformation of the outcome variable in Table B3.

³⁴[Griffin et al. \(2020\)](#) conduct a simulation exercise that considers the statistical performance of difference-in-differences research designs when estimating the impacts of state-level opioid policies on overdose mortality. They find that many commonly-used DiD study designs and model specifications have low statistical power to detect an effect, as well as high rates of Type I error. It is possible that one cause of these findings is related to bias in the TWFE estimator, which can induce fragility of the estimated effects to choice of model.

The number of lives saved per year can be roughly approximated using a back-of-the-envelope calculation: a 10.8% reduction in 2,267 counties with a mean of 0.62 opioid deaths per quarter implies 607 deaths averted per year in my 38-state sample, or around 8,000 deaths during the 2001-2013 period. A counterfactual simulation exercise using the dynamic Poisson specification reported in Table B3, Column (4), where deaths are predicted as if every state had implemented a PDMP in Quarter 1 of 2004, produces a similar number, around 9,000 deaths fewer than the 73,000 deaths that occurred in the 38 states in the study during that period.

Finally, I account for emerging concerns that cracking down on prescription opioids has led to substitution towards more-dangerous heroin use, which might offset the gains documented in overdose mortality linked to prescription opioids, by estimating the impact of PDMP introduction on all opioid and heroin-linked deaths, indicated by multiple causes of death codes T40.1-T40.4 and T40.6. I find no meaningful indicators that the main effect is attenuated by substitution towards these other drug categories, and if anything considering all types of opioid overdose death increases the implied effect. The results are reported in Table 4, Column (2). While precision is reduced in this specification, an 8.2% reduction total opioid overdose deaths off a mean of 1.37 opioid deaths per quarter implies a larger net number of deaths averted, around 13,000; similarly, in the counterfactual Poisson simulation, about 11,000 deaths are averted out of 161,000 during that period.³⁵

5.3 Opioid abuse declines

Additionally, I construct an indicator of rates of opioid use disorder using the MarketScan medical claims data, described above, and use the indicator for a separate measure of the impact of PDMP introduction on opioid harms. Considering opioid use disorder and opioid abuse provides additional insight, beyond fatalities, on whether PDMPs had an impact on the overall burden of disease arising from the opioid epidemic. Additionally, in Section 7.1 I will use the fact that this indicator of abuse is linked to the medical claims and work data to make statements about the private costs and benefits of reducing access to prescribing, by measuring the impact of PDMP introduction on rates of opioid abuse in chronic pain patients.

I identify episodes of opioid use disorder using opioid dependence and opioid overdose diagnosis codes, as well as observed prescriptions for OUD treatment medications, and claims in mental health and substance abuse facilities, and use a custom episode grouper to link related claims across time.³⁶

³⁵In prior versions of this paper, a temporary 1-year increase in heroin overdose deaths was reported, followed by a subsequent decline. This result is especially fragile to choice of specification, as detailed above, and in particular is not robust to estimation using the procedure in [de Chaisemartin and D'Haultfoeuille \(2020\)](#), and so is not included in subsequent revisions.

³⁶Specifically, I consider a person to have developed opioid use disorder if they meet two out of three criteria in a given month: (a) having any ICD-0 code for opioid dependence or overdose (one of 304.00-304.03, 304.70-304.73, 305.50-305.53, 965.00-965.02, 965.09, E850.0-E850.2), (b) having a opioid use disorder treatment prescription, or any claim containing

I find low rates of opioid use disorder – about 1.1% of enrollees prescribed oxycodone were later observed with a claim relating to opioid use disorder, and only about 23,000 enrollees in the full dataset are identified as having opioid use disorder. These numbers are in line with findings of other research on opioid use disorder in claims data (Barocas et al., 2018), and are likely to represent an undercount of the true rates of opioid abuse in this population. This undercount reflects widespread under-diagnosis and under-treatment of opioid use disorder in general.³⁷

I estimate the following models for the development of opioid use disorder, using logistic regression for my primary specifications:

$$1(ODD)_{ist} = \alpha + \gamma_e + \lambda_t + \sum_{\tau} \sigma_{\tau} D_{\tau,st} + \beta X_{it} + \epsilon_{iest} \quad (5)$$

$$1(ODD)_{ist} = \alpha + \gamma_e + \lambda_t + \beta_1 1(PDMP_{st}) + \beta_2 X_{it} + \epsilon_{iest} \quad (6)$$

where i is individual, s is state, e is entry-cohort, and t is quarter, and $1(ODD)_{ist}$ is a binary outcome variable indicating whether the enrollee meets the criteria for opioid use disorder.

The results are depicted in Table 5, column (1), and Figure 5. The point estimates can be interpreted as an approximation of the reduction in the relative risk for small numbers at this order of magnitude, and column 1 thus reports an approximate 6% decline in cases after the introduction of a PDMP.³⁸ The event study shows a pattern that is similar to the event study for opioid overdose deaths. The reduction in opioid abuse episodes phases in over the first two years and persists for at least four years.

6 Empirical Analysis: Consequences of Reduced Opioid Prescribing

Next, I consider the consequences of reduced opioid prescribing on the treatment of pain. Pain management is a classic ‘grey area of medicine’ – the evidence base is mixed, and patient preferences should be important to determining a course of treatment (Chandra, Cutler and Song, 2012). Although opioid therapy has become controversial in recent years, and the evidence base is very incon-

buprenorphine as an active ingredient, or (c) a claim in a mental health or substance abuse facility, or with a mental health or substance abuse provider.

³⁷The rate of 3 per 1,000 implies approximately 941,000 people with opioid use disorder in the U.S. during this time period, this is about half the estimated 1.8 million Americans with OUD. However, a 50% undercount is likely the lower bound, as it is likely that people suffering from opioid use disorder are disproportionately uninsured or on Medicaid.

³⁸These results are very similar to results estimated with a linear probability model using individual or entry-cohort fixed effects, but the logistic regression results are presented because they are more interpretable, especially when comparing across groups with different rates of opioid use disorder. For example, for the full sample in Table 5, Column (1), the percentage point reduction estimated by the linear probability model is -0.0001415; with a mean rate in the sample of 0.00156, 100,000 individual-quarters would have 156 cases of opioid use disorder, and the estimated effect of -0.0000875 would represent 8.75 fewer cases after the introduction of a PDMP, or a 5.6% reduction in cases, which is very similar to the relative effect estimated with the logistic regression.

clusive on safety and efficacy, as is explored in Sections 1 and 2, other approaches to pain management, especially interventional pain management (e.g., epidural steroid injections), and surgery are also controversial, as evidence on their effectiveness has been mixed.³⁹ Heavy utilization of opioid therapy to manage pain, rather than physical therapy, interventional pain management approaches, or surgery, likely found favor in part due to the fact that prescribing opioids is far cheaper for the payer and less time-intensive for the provider.⁴⁰ But if there are agency issues between the insurer and patient, or doctor and patient, reducing opioid availability and shifting towards more expensive therapies which may treat the root of the pain problem rather than the symptom may be preferable for some patients, especially if they consider their private risks of addiction or overdose (Thomas et al., 2015).

I exploit rich claims data on healthcare utilization (the Marketscan databases) to uncover how medical treatment of pain changes in response to the introduction of a PDMP. For these datasets, I aggregate raw claim-level data to observations at the individual-quarter as described in Section 3.2, and estimate:

$$Y_{it} = \alpha + \gamma_i + \lambda_t + \sum_{\tau} \sigma_{\tau} D_{\tau,ist} + \beta X_{it} + \epsilon_{it} \quad (7)$$

$$Y_{it} = \alpha + \gamma_i + \lambda_t + \beta_1 1(PDMP_{ist}) + \beta_2 X_{it} + \epsilon_{it} \quad (8)$$

where i is individual, s is state, and t is quarter, and Y_{it} is an outcome variable constructed from Marketscan raw data such as $\ln(\text{morphine milligram equivalent [MME] per quarter} + 1)$ or $\ln(\text{inpatient and outpatient spending} + 1)$. For binary outcome variable specifications, I estimate both a linear probability model and a logistic model. For count data (days absent under workers compensation and short-term disability), I estimate a Poisson pseudo-maximum likelihood model.

6.1 Utilization of opioids for the treatment of pain

I first document that the overall effect of PDMP introduction on opioid prescribing for the Marketscan medical claims dataset, which represents the experience of approximately 175 million American individuals with employer-sponsored health insurance, comports well with the results from the DEA ARCOS controlled substance monitoring dataset, which in theory contains a complete accounting of all opioids that are distributed. The results of estimating Equations 7 and 8 are presented in Figure 6 and Table 6.

The results show a statistically significant decline in Morphine Milligram Equivalent (MME) of

³⁹As noted in a previous footnote, in 2014 Medicare cut reimbursement rates to many kinds of interventional pain management for this reason. Certain kinds of surgery, such as spinal fusion surgery for back pain, are also considered questionably evidence-based (Christensen, 2004).

⁴⁰Note that other developed countries, particularly in Europe, favor an integrative approach to pain management, and use far fewer opioids per capita.

Schedule II opioid prescribed for all individuals in the sample, post PDMP introduction. I present results for $\ln(\text{MME} + 1)$ as the dependent variable in Column 1.⁴¹ The point estimate from Panel 1 indicates that there is a 0.5% decrease in the *geometric* mean of the distribution of $\text{MME} + 1$ for the full sample.⁴² Column 4 considers the extensive margin of prescribing using a binary outcome variable. The results indicate that the introduction of a PDMP reduces the overall probability of an enrollee receiving opioids in a quarter by 0.0562 percentage points, i.e., from about 2.47 percent overall, to about 2.41 percent. This estimate captures reductions on the extensive margin, i.e., patients entering or exiting opioid therapy, and (unlike the specification in Column 1) does not capture changes in the level of prescribing, and specifically would not capture dosage reductions for high-utilization chronic pain patients. Column (5) similarly reports the extensive margin effects, estimated with a logistic regression. As discussed previously, the point estimate can be interpreted as the reduction in relative risk, and it closely matches the estimate of Column (4), which implies a 2.3% reduction in overall enrollees receiving opioids in a given quarter.⁴³ Finally, considering the results in Columns (2) and (3) assists in understanding the change in overall prescribing levels by using the raw level of MME dispensed in the quarter as the dependent variable. Column (2) reports a Poisson pseudo-maximum likelihood specification implying a reduction in overall prescribing quantity of around 9%. Finally, the raw MME per quarter OLS specification gives an unbiased estimate of the average reduction in per capita opioids dispensed per quarter in the Marketscan claims data, and thus is useful for comparison to the ARCOS result. The result demonstrates that the introduction of a PDMP is associated with an average of 6.43 MME fewer Schedule II opioids prescribed per patient per quarter (from a mean of 60.99 MME, so a decrease of about 10.5% in the total quantity). As discussed in Section 3.2, the average MME per capita in the ARCOS dataset is somewhat higher than in the Marketscan claims data (106.7 MME versus 60.99 MME per capita per quarter) likely because Marketscan does not include the elderly or the poor, and does not capture most illegal or high-volume suspicious diversion. The point estimate of the reduction in MME per capita for the Marketscan prescription claims dataset represents a 10.5% decline, compared to an 11.2% decline in MME per capita for the ARCOS controlled substances monitoring data.

That these reductions are so similar across the two datasets is suggestive that PDMPs are af-

⁴¹As detailed above in Section 3.2, due to the fact that individuals tolerate physically to opioids very quickly, the opioid distribution for the overall Marketscan sample is highly skewed, with a large number of zeros as well as a very long right tail. I utilize $\ln(\text{MME} + 1)$ as the outcome variable in my baseline specification because it is the most straightforward transformation with which to deal with outliers, while also retaining the large number of zeros in the sample and information about the quantity/level consumed in quarter. I consider an alternate transformation, the inverse hyperbolic sine, in Table B4; results are similarly robust to other concave transformations such as $\sqrt{\text{MME}}$ or $\ln(\text{MME} + 0.01)$, not shown.

⁴² $(e^{\beta_1} - 1) * 100$ is the percent change in the geometric mean of $\text{MME} + 1$ associated with going from no PDMP to PDMP. The geometric mean captures the central tendency of the distribution and is in this case very different from the arithmetic mean – 60.99 MME per quarter (arithmetic) versus 0.15 MME per quarter (geometric).

⁴³Using the point estimate and mean from Column (4), $1 - (.0247 - .000562)/.02470 = .0228$, which is very similar to the estimate in column (5).

fecting prescribing and dispensing roughly proportionally for disparate U.S. populations, including those with employer-sponsored insurance. It also demonstrates that the reduction in overall opioid availability is not simply due to a reduction in high-volume diversion through pill mills or pharmacy theft, and is occurring for employer insured patients as well. The next section considers the characteristics of medical users who experience the sharpest declines in opioid access.

6.1.1 Margins along which opioid prescribing is reduced

As described in Section 2, doctors and public health practitioners are increasingly questioning the role of opioids in the treatment of long-term, chronic non-cancer pain. Recent legal initiatives have also targeted prescribing for acute conditions, based on the concern that excess prescribing in the acute setting can lead to abuse among those patients, or that left-over pills might be diverted (Miller, 2015).

I attempt to better understand the margins along which the PDMP-induced reduction in opioid prescribing is occurring, using several approaches. First, as detailed in Section 3.2, I group opioid prescription claims into ‘opioid therapy episodes’ based on temporal proximity and generic ingredient, allowing me to observe whether opioid therapy is *initiated* in a given quarter when the enrollee is not currently taking an opioid. Recall from Section 3 that most prescriptions for opioids (around 80%) end after one prescription. Table 7, Panel A, demonstrates that there is a relative reduction in whether opioid therapy is initiated at all of about 2%. Given that the majority of opioid therapy episodes are short-term therapy for acute pain conditions such as broken bones, it is likely that this reduction in initiation of opioid therapy represents a reduction in the treatment of acute pain episodes using opioid therapy.

In Table 7, Panel B, I disaggregate the distribution of enrollee-quarter observations for MME opioid into bins (where thresholds are set according to CDC prescribing guideline thresholds discussed in Footnote 2 and explained in the table notes), and define the outcome variables as binary (0/1) for whether or not a quantity of opioid within the thresholds for that bin were observed in a given quarter. This exercise demonstrates that there is shrinkage at the highest parts of the distribution, representing continuous chronic prescriptions (Panel B-6, representing patients receiving approximately 8,100 MME opioid in the quarter or greater, or 90 MME per day continuously, and Panel B-5, between 50 MME and 90 MME per day). It is unclear from this exercise whether Panels B-2 through B-4 see a decline; the insignificant results for the lower bins are consistent with a story where individuals in the top bins have been moved to lower bins, while those in the lower bins have been discontinued. But it is also consistent with a story where individuals move from the top bins to “no opioids;” this would be a story where PDMPs primarily function to catch high-volume “doctor shoppers,” while the rest of the distribution is unaffected. The significant reduction from the top bins suggests that people on

chronic opioid therapy are receiving a reduction in dose, or are being discontinued entirely, and the result is of large magnitude, representing an 11 percent reduction in the number of patients in that bin.

Taking the results from Panels A and B together, I conclude that the reduction in opioid prescribing is occurring at all parts of the prescribing distribution – for initial prescriptions (most of which will be for acute pain), and for existing high-volume users (likely, individuals with chronic pain, and/or doctor-shoppers).

6.1.2 Robustness and specification checks

Robustness checks presented in Table B4 indicate that the main Marketscan results for prescribing are robust to a variety of other alternate specifications, including state-specific linear time trends and most alternate legal codings, as well as alternate concave transformations of the dependent variable. Additionally, I consider in Appendix D concerns related to the discussion in Section 3 regarding the manner in which the Marketscan sample is constructed. Because Marketscan data is not a representative national sample, but rather is a sample obtained from employers (meaning that individuals working for the same firm, whose characteristics are likely to be correlated, enter or leave the sample in groups as their employers join or exit), unobserved heterogeneity could be introduced as different types of individuals move into and out of sample in blocks. In Appendix Table D1, I consider robustness of the baseline specification to alternate fixed effects models that account for this sample construction, and in Appendix Figure D1 I show that entry-cohort fixed effects (which are much lower-dimensional) are sufficient to control for unobserved heterogeneity in cohorts.

6.2 Welfare effects of reduced opioid prescribing on patients with chronic pain

As established in Section 6.1, opioid prescribing declines for both initial and high-quantity chronic prescriptions, and thus likely affects patients with both acute and chronic pain. If a reduction in opioid prescribing falls on chronic pain patients, this would raise potential welfare concerns arising from burdens associated with their need to change their approach to pain care. Specifically, we might expect to observe a compensating increase in other kinds of medical spending to manage their pain, and might expect to observe other negative outcomes such as reduced labor supply.

I exploit the richness of the Marketscan data to train a machine learning model (described in Appendix E) that can identify a cohort of chronic pain patients for study using an array of outcomes on their health, labor supply, and opioid use disorder. Appendix Table E2 depicts demographic details about this selected cohort. Compared to the full sample (excluding the chronic pain cohort), chronic pain patients are more likely to be older, female, and working in jobs that appear likely to be blue-

collar (unionized and paid hourly instead of salaried, and employed in sectors like manufacturing rather than sectors like finance, insurance, and real estate). As can be seen in Table 8, this cohort is indeed prescribed a much higher volume of opioids – about five times as much volume as the full sample (318 MME per individual per quarter compared to 60.99 MME per individual per quarter), and, like the full sample, they similarly experience a reduction in total opioid prescribing after PDMP introduction. Reflecting overall worse health, total medical spending in this cohort is three times that of the full sample, and they miss four-and-a-half times more work under workers compensation and short term disability than the full sample.

6.2.1 Substitution towards more expensive inpatient and outpatient spending

In Table 8, I display estimates of the impact of a PDMP on overall inpatient and outpatient spending on the chronic pain cohort, and compare to the full sample of enrollees. The event study for the impact of PDMPs on inpatient and outpatient spending for this chronic pain sample is depicted in Figure 7. Post-PDMP, there is a discontinuous jump in medical spending that is persistent and ranges from approximately 5% initially to 10% in later years.

Next, I use the richness of the claims data to segment and characterize the chronic pain cohort’s health episodes using an unsupervised learning natural language processing algorithm, latent Dirichlet allocation, which classifies data into topics where each observation in a collection is substantially similar. Specifically, the algorithm reads in all the rich details provided in the claims data related to a temporally-linked episode of care – all diagnosis codes, procedure codes, visits to different types of physicians, etc. – and groups, in an unsupervised manner, these episodes into similar kinds of experiences where words (in this case, medical codes, etc.) appear together frequently. By identifying interpretable groupings of types of care experienced by chronic pain patients, the approach gives an enriched qualitative understanding of the experiences of the chronic pain cohort.

A descriptive listing of the 15 topic categories is provided in Appendix Table F1, and several example topic groups are presented in Appendix F.⁴⁴ Topic 2, depicted in Figure F1 identifies claims related to inpatient surgeries and hospitalizations related to musculoskeletal pain and joint disorders, while Topic 8, depicted in Figure F2 identifies claims related to outpatient treatment of pain conditions via physical therapy, chiropractic medicine, etc.⁴⁵ In Table F1 and F2, I group these two types of medical claims episodes together as “pain episodes.” In Table F2, I split the chronic pain cohort’s sample over time according to the type of medical episode, and show specifically that total

⁴⁴All fifteen groups are presented in an interactive manner on https://angelakilby.com/html/LDA-pain_patients.html.

⁴⁵The algorithm also illustrates that the chronic pain cohort experiences other types of medical episodes. For example, the algorithm identifies ER visits as a category of medical episode (Topic 12, depicted in Figure F3), routine care like office visits for preventative medicine and routine tests (Topic 4, depicted in Figure F4), and other serious types of medical experience like cancer (Topic 6, depicted in Figure F5). Note that the topic ordering in the LDA topic model is not meaningful.

spending on pain-related episodes grows more costly after the introduction of a PDMP and associated reductions in opioid prescribing. This substitution towards more expensive kinds of pain care is a financially costly result to the insurer of prescribing restrictions. As discussed above in the introduction to this section, this substitution away from opioids and towards alternate kinds of pain care delivered in inpatient and outpatient settings may be welfare-decreasing or increasing, depending on patient preferences over pain management inputs, many of which have a questionable evidence base, and whether there were agency problems between individuals and their insurer in getting more expensive procedures such as surgery approved. The substitution is suggestive of an increased burden for pain patients, motivating examination of work productivity effects for these patients in the next section.⁴⁶

These results, in combination with the results in Section 5.1, demonstrate that the reduction in opioid prescribing is occurring not only for non-medical users, people engaging in diversion, or pill mills, but also for medical users with legitimate needs for pain management. Because opioid medications are cheap, these results show that overall medical costs for chronic pain patients will *increase* with a reduction in opioid availability. Given past medical studies, detailed in Section 2, have indicated that higher opioid use is associated with higher medical spending, this result contextualizes those findings, indicating the observed associations in those studies are likely not causal. Nonetheless, it is not clear from this result alone that patients are being harmed by an opioid crackdown from a private welfare perspective, given insured patients do not bear the majority of the cost of this increase in spending. The substitution towards new and more expensive kinds of pain management care is suggestive, but not definitive, that their pain may not be properly managed after the reduction in opioid availability.

6.2.2 Increased absences and reduced labor supply

I exploit a unique feature of the Marketscan claims dataset – its linkage to workers’ compensation claims and short-term disability data – to examine the work ramifications of reducing opioids for workers that are likely to need them to manage pain. As discussed above, the chronic pain sample misses four-and-a-half times more work under workers compensation and short term disability than the full sample. Workers compensation specifically is often cited as a domain where opioids are overprescribed: about 70% of workers missing days of work under workers’ compensation received

⁴⁶In other analysis (not shown), I consider substitution to alternate forms of pain management in an alternate manner, by disaggregating spending by hand according to the associated diagnosis and procedure codes, using rough categories of procedures that might provide pain management as a substitute for opioid therapy. For example, under interventional pain management, I included procedures such as epidural steroid injections and spinal cord stimulation; under surgery I included procedures such as spinal fusion surgery and joint replacement. I found notable substitution, e.g., towards spending on surgery – a 1.3% increase. However, claims data is highly disaggregated and it can be difficult to identify all expenses associated with a given episode of surgery care; the above approach groups episodes of care in a natural manner and characterizes them according to the most common types of procedures and diagnoses in each episode.

an opioid prescription. There has also been increasing alarm among workers' compensation payers that narcotics use is contributing to increases in total medical and indemnity (lost wages) costs for injured workers. A typical study concluded that "workers prescribed even one opioid had average total claims costs 4-8 times greater than claimants with similar claims who didn't get opioids... the cost of a workplace injury is nine times higher when a strong narcotic like OxyContin is used than when a narcotic is not used (Rosenblum, 2012)." As discussed in the introduction, these studies usually use matched controls, (see e.g., Webster, Verma and Gatchel (2007); Franklin et al. (2008)) and are plagued by endogeneity problems. For each chronic pain enrollee, I consider all absenteeism across workers compensation and short-term disability. As with workers' compensation, the literature on the role of opioids in short-term disability has been observational, and has typically found that medical opioid use is associated with *decreased* work and function (Ashworth et al., 2013; Mahmud et al., 2000; Franklin et al., 2008).

In the sample of chronic pain patients, I document a strong and persistent increase in days missed under workers' compensation and short-term disability that coincides with the introduction of a Prescription Drug Monitoring Program. The point estimates are displayed in Table 8 and the equivalent event study is depicted in Figure 8; the introduction of a PDMP coincides with a persistent increase in days missed of about 6.2% on average. The change in days missed seems to occur on the intensive margin; as is shown in Table B5, the increase does not come from an increase in the frequency of a worker being injured or taking short term disability, nor does the introduction of a PDMP seem to affect whether a workers' compensation injury results in having some days lost versus none. Table F2 shows that like the increases in inpatient and outpatient spending, these increases in missed days of work are associated temporally with the treatment of pain conditions specifically.

The decline in work output for chronic pain patients on workers' compensation and disability after an opioid crackdown is important because it is a clear demonstration of direct welfare losses due to the reduction in availability of opioids and subsequent increase in pain. Aside from being of direct policy relevance to workers' compensation payers and policymakers due to increased costs of lost wages, the results linking reduced opioid use to increased days missed provide evidence that opioid use is causally linked to *improved* ability to function in the long term for people needing pain management, a result with broader implications for individuals in pain. As discussed above, prior evidence on long term opioid use and overall function has been limited to observational studies because no randomized trials of opioids have been conducted that last more than one year; this paper's quasi-experimental approach uncovers a result that is of the opposite sign to the observational approach.

In conjunction with the results on increased inpatient and outpatient spending in Section 6.2.1, these results indicate that pain and function might be worsening for pain patients after an opioid

crackdown, suggesting important welfare losses for individuals with pain. Losing access to opioid pain management may reduce chronic pain patients' work output, and harm overall functionality, freedom from suffering, life satisfaction, and happiness.

7 Empirical Analysis: Quantifying Welfare Tradeoffs of Reducing Opioid Use

I have documented four main effects of introducing a Prescription Drug Monitoring Program. First, opioid prescribing declines broadly, affecting both acute and chronic prescribing. Second, opioid-related harms are reduced, specifically, rates of opioid use disorder and opioid overdose deaths. Third, substitute spending on alternative (and more expensive) inpatient and outpatient medical care increases. Finally, days missed for short-term disabled and injured workers increase.

7.1 Distributional concerns

First, I consider what my estimates suggest about the *private* welfare tradeoffs facing pain patients. As discussed above, there is controversy in the medical and public health communities over how to best understand opioid abuse and opioid harms. Some practitioners believe that opioids can be used safely by individuals with legitimate medical need, and opioid abuse is primarily a result of diversion and theft from the legitimate medical care system. However, an emerging alternative viewpoint is that widespread legitimate medical opioid use is the primary driver of opioid abuse, addiction, overdose, and death – that the medical care system is producing “iatrogenic addiction” and is thus largely responsible for opioid-related harms.

If the first story dominates, there would be a tradeoff in the incidence of costs and benefits of reduced overall access to opioids, from the perspective of social welfare. If we reduce access to opioids to people in need, their welfare may go down as they are unable to work and as they substitute towards alternate, possibly less effective, pain management. The benefits of regulation accrue to the individuals who are less able to access diverted drugs, and consequently are less likely to experience addiction and death. If the second story dominates, pain patients are experiencing both the costs of regulation, in terms of lost work and functionality, and the benefits, in that they themselves are less likely to be come addicted, overdose, and die.

An additional nuance comes in distinguishing the welfare tradeoffs of acute prescribing, which may be different than prescribing for chronic pain. While acute prescribing has typically been considered less likely to lead to the development of opioid use disorder than long-term opioid therapy, it is also potentially less beneficial from a welfare perspective than prescribing for long-term chronic

pain.

To address this, I first directly consider the impact of the reduction in opioid prescribing for the group who has the most documented harms: the chronic pain cohort. In Table 5, columns (2) and (3), I consider the rate of opioid use disorder episodes in the full sample excluding the chronic pain sample, compared to the chronic pain sample. Strikingly, despite much higher baseline risks of opioid use disorder in the chronic pain sample, the relative reduction in opioid use disorder risk is more than double in magnitude for the non-chronic pain cohort than the chronic pain cohort, and the point estimate for the chronic pain sample is not distinguishable from zero. As reported in Table 10 (and considered alongside Table 8), both groups saw declines in both probability of receiving a prescription and in overall prescribed volume. The non-chronic pain sample could be benefiting from reduced acute prescribing or from reduced diversion of opioids that were not prescribed to them, either of which might reduce rates of opioid abuse in this group. In contrast, the benefits to the chronic pain sample specifically are decidedly unclear, especially in context of the apparent costs in terms of increased pain and loss of work productivity. Benefits in the form of reductions in opioid use disorder appear to be largely accruing to the non-chronic pain cohort, while costs appear to mostly accrue to the chronic pain cohort.

In Tables 10 and 11, I provide additional suggestive evidence. I disaggregate prescribing results and overdose deaths according to demographic variables, including sex, age group, urbanicity, and available class indicators. (I utilize coarse employment variables in MarketScan, including union/non-union and salaried/hourly, and education variables in the CDC death data.) While many of the estimates are noisy, especially for overdose deaths, prescribing restrictions appear to hit groups that are older, geographically diverse across urban and rural areas, and more male, whereas opioid overdose deaths are averted in groups that are relatively younger, more urban, and more female. While class indicators are not directly comparable, prescribing restrictions and overdose deaths averted appear to both be concentrated in groups with lower socioeconomic status markers: blue-collar workers in the prescribing data and less-educated groups in the overdose deaths data.

Taken together these results suggest that the reduction in opioid prescribing brought about by heightened opioid regulations has had important social benefits, but has also exerted disparate impacts on certain groups, namely individuals suffering from chronic pain, who bear most of the costs and accrue relatively few benefits.

7.2 Overall welfare

Finally, I use my results on prescribing, opioid abuse and overdoses, substitute medical spending, and work output to undertake a broad social welfare calculation, attempting to integrate the results to broadly quantify the costs and benefits of a blanket reduction in opioid prescribing. First, in the

Marketscan claims data, I note that most of the costs and benefits to society of a reduction in opioid prescribing are directly monetized. For the 340,660 patients in the pain cohort, I sum spending on prescriptions (which includes the decrease in spending on opioids, as well as any changes in spending for other drugs) and inpatient and outpatient utilization, thus encompassing all observed medical spending for the employer-sponsored insurance. I then add lost wages (through an imputed daily wage based on the total indemnity payment for a claim), and medical payments under workers' compensation, and thus consider the change in total spending on healthcare and lost wages, post-PDMP. The result of this exercise is found in Table 9. There is a 4.5% increase in total costs, and the mean total cost for this group is \$3,920 per quarter, so this increase represents about \$177 in increased spending per quarter for these enrollees. Assuming the Marketscan dataset represents the population of employees with employer-provided health insurance as discussed in Section 3, these 340,660 individuals represent roughly 11.6 million people nationally, and an increase in spending of \$8.2 billion annually for their health care and lost wages.⁴⁷

These costs are traded off against the value of an 8.2% overall reduction in opioid and heroin overdose deaths; as discussed in Section 5.2, or 8,000 to 13,000 deaths averted over 13 years. To make a rough comparison, I employ a figure for the Value of a Statistical Life (VSL) of \$7.4 million; this implies the benefits in lives saved according to this metric are approximately \$4.6-\$7.4 billion. Further, in order to capture reductions in other costs to society of substance abuse and addiction, I use estimates from other literature on rates of treatment admissions, emergency department visits and incarcerations that are benchmarked against opioid overdose death rates as a proxy for levels of underlying abuse. These estimates suggest that one overdose death represents approximately 10 treatment admissions and 32 emergency department visits, and incarceration costs per opioid user are estimated to be approximately 3 times hospital and ED costs ([Center for Integrated Behavioral Health Policy, 2011](#); [Kassed, Levit and Hambrick, 2007](#); [HCCI, 2013](#)). As such, my estimates of 615-1000 deaths averted per year suggest 6,150-10,000 treatment admissions were prevented at an average cost of \$7,230 per admission. This adds up to \$44-72 million. Similarly, there are approximately 32 fewer emergency department visits for every overdose death, so 20,000-32,000 ED visits were prevented at an average cost per ED visit of \$1,200. This figure totals to \$24-38 million. Finally, incarceration costs per user, estimated at approximately 3 times hospital and ED costs, total to \$72-115 million. In total, these additional considerations add \$140-\$225 million, and the total benefits from reduced opioid addiction, overdose, and death can thus be roughly valued at \$4.8-7.6 billion.

⁴⁷Note that the last line of Table 9 contains the implied increase in costs if I were to consider each subsample independently. Though the point estimate is noisily estimated due to the large numbers of individuals who have inpatient and outpatient spending that are not affected by a reduction in opioids, the estimate of \$10.9 billion for all 4,944,255 enrollees in Column (1) is reasonably close to the estimate for the combined sample enrollees in Column (2), indicating that my chronic pain sample may mostly if not entirely capture the individuals in the Marketscan dataset who have costly needs for pain management.

This is a smaller figure, but on the same order of magnitude, as the \$8.2-\$10.9 billion figure for costs derived above.

The above two numbers are subject to a number of major limitations. Costs may be underestimated for many reasons. First, I calculate costs of restricting opioids based on a subpopulation that is most in need of pain management; I exclude the lost work of people in acute (temporary) pain, and all other individuals who need pain management but do not fall into my group of predictive enrollees. Second, I do not measure lost work output for any individuals other than those injured in workers' compensation or on short-term disability. Third, I cannot observe the direct psychological costs of increased pain. As explored in [Case and Deaton \(2015a\)](#) and [Case and Deaton \(2015b\)](#), increasing pain among certain subsets of the US population, especially middle-aged white men, is associated with increased distress, declining mental health, and increased mortality, including suicide.

Benefits of increased regulation may also be underestimated. Costs to society of drug abuse, addiction, overdose, and death accrue in many forms, such as increased need for family services. There are also non-monetary costs in the form of suffering by individuals with opioid use disorder, their families, and their community.

8 Conclusion

This study exploits a natural policy experiment, state-level variation in the timing of the introduction of Prescription Drug Monitoring Programs, to examine the costs and benefits of reducing access to opioids. The findings of this study are several. First, PDMPs, an increasingly important policy lever for reducing access to opioids, do indeed reduce the use of opioids, and achieve major policy successes in reducing opioid abuse and opioid overdose deaths. However, they have had several unintended consequences. Pain management is altered – as access to prescription opioids is curtailed, chronic pain patients substitute towards alternate and more-costly forms of medical pain management. Functionality and ability to work decline: days absent for injured and disabled workers increase, creating clear welfare losses for those individuals, and suggesting that other unobserved welfare losses attributable to losing access to pain management (in terms of functionality and overall life satisfaction) are also likely.

Back-of-the-envelope calculations suggest that the regulation of prescription opioids involves costs and benefits that are both large: welfare losses from a crackdown on opioids are on the order of \$8.2-\$10.9 billion in increased medical spending and lost wages, whereas welfare gains are on the order of \$4.8-\$7.6 billion from reduced opioid use disorder, overdose, and death.

References

- AAACN. 2001. "Managing Patients' Pain." *AAACN Viewpoint Special Issue: Treating Pain*, 23(1).
- Ailes, Elizabeth, April Dawson, Jennifer Lind, Suzanne Gilboa, Meghan Frey, Cheryl Broussard, and Margaret Honein. 2015. "Opioid Prescription Claims Among Women of Reproductive Age—United States, 2008-2012." *MMWR. Morbidity and Mortality Weekly Report*, 64(2): 37–41.
- Aitken, Murray. 2013. *Declining Medicine Use and Costs: For Better or Worse?* IMS Institute for Healthcare Informatics.
- Aitken, Murray. 2019. "Medicine Use and Spending in the U.S.: A Review of 2018 and Outlook to 2023." IQVIA Institute for Human Data Science.
- Alford, Daniel P. 2016. "Opioid Prescribing for Chronic Pain – Achieving the Right Balance through Education." *New England Journal of Medicine*, 374(4): 301–303.
- Alpert, Abby E., Sarah E. Dykstra, and Mireille Jacobson. 2020. "How Do Prescription Drug Monitoring Programs Reduce Opioid Prescribing? The Role of Hassle Costs versus Information." National Bureau of Economic Research, Inc 27584.
- Alpert, Abby, William Evans, Ethan M.J. Lieber, and David Powell. 2019. "Origins of the Opioid Crisis and Its Enduring Impacts." National Bureau of Economic Research w26500, Cambridge, MA.
- Anson, Pat. 2015. "Chronic Pain Groups Blast CDC for Opioid Guidelines." Pain News Network.
- Ashworth, Julie, Daniel J. Green, Kate M. Dunn, and Kelvin P. Jordan. 2013. "Opioid Use Among Low Back Pain Patients in Primary Care: Is Opioid Prescription Associated with Disability at 6-Month Follow-Up?" *Pain*, 154(7): 1038–1044.
- Ault, Alicia. 2015. "CDC Opioid Guideline Lands, Controversy Continues." Medscape.
- Ballantyne, Jane C. 2006. "Opioids for Chronic Pain: Taking Stock." *Pain*, 125(1): 3–4.
- Ballantyne, Jane C. 2012. "'Safe and Effective When Used as Directed': The Case of Chronic Use of Opioid Analgesics." *Journal of Medical Toxicology: Official Journal of the American College of Medical Toxicology*, 8(4): 417–423.
- Ballantyne, Jane C., and Naomi S. Shin. 2008. "Efficacy of Opioids for Chronic Pain: A Review of the Evidence." *The Clinical Journal of Pain*, 24(6): 469–478.
- Bao, Yuhua, Katherine Wen, Phyllis Johnson, Philip J. Jeng, Zachary F. Meisel, and Bruce R. Schackman. 2018. "Assessing The Impact Of State Policies For Prescription Drug Monitoring Programs On High-Risk Opioid Prescriptions." *Health Affairs*, 37(10): 1596–1604.
- Bao, Yuhua, Yijun Pan, Aryn Taylor, Sharmini Radakrishnan, Feijun Luo, Harold Alan Pincus, and Bruce R. Schackman. 2016. "Prescription Drug Monitoring Programs Are Associated With Sustained Reductions In Opioid Prescribing By Physicians." *Health Affairs*, 35(6): 1045–1051.
- Barocas, Joshua A., Laura F. White, Jianing Wang, Alexander Y. Walley, Marc R. LaRochelle, Dana Bernson, Thomas Land, Jake R. Morgan, Jeffrey H. Samet, and Benjamin P. Linas. 2018. "Estimated Prevalence of Opioid Use Disorder in Massachusetts, 2011–2015: A Capture–Recapture Analysis." *American Journal of Public Health*, 108(12): 1675–1681.

- Berndt, E. R., H. L. Bailit, M. B. Keller, J. C. Verner, and S. N. Finkelstein.** 2000. "Health Care Use and At-Work Productivity Among Employees with Mental Disorders." *Health Affairs (Project Hope)*, 19(4): 244–256.
- Bernstein, Lenny.** 2015. "CDC Urge Doctors to Curb Opioid Prescriptions: "The Risks are Addiction and Death, and the Benefits are Unproven"." The Washington Post.
- Blei, David M., Andrew Y. Ng, and Michael I. Jordan.** 2003. "Latent Dirichlet Allocation." *The Journal of Machine Learning Research*, 3(null): 993–1022.
- Board of Pharmacy.** 2012. "New Mexico Regulation and Licensing Department."
- Borusyak, Kirill, and Xavier Jaravel.** 2016. "Revisiting Event Study Designs."
- Braden, Jennifer B., Alicia Young, Mark D. Sullivan, Brian Walitt, Andrea Z. LaCroix, and Lisa Martin.** 2012. "Predictors of Change in Pain and Physical Functioning Among Post-Menopausal Women with Recurrent Pain Conditions in the Women's Health Initiative Observational Cohort." *The Journal of Pain*, 13(1): 64–72.
- Buchmueller, Thomas C., and Colleen Carey.** 2018. "The Effect of Prescription Drug Monitoring Programs on Opioid Utilization in Medicare." *American Economic Journal: Economic Policy*, 10(1): 77–112.
- Busse, Jason W., Li Wang, Mostafa Kamaleldin, Samantha Craigie, John J. Riva, Luis Montoya, Sohail M. Mulla, Luciane C. Lopes, Nicole Vogel, Eric Chen, Karin Kirmayr, Kyle De Oliveira, Lori Olivieri, Alka Kaushal, Luis E. Chaparro, Inna Oyberman, Arnav Agarwal, Rachel Couban, Ludwig Tsoi, Tommy Lam, Per Olav Vandvik, Sandy Hsu, Malgorzata M. Bala, Stefan Schandelmaier, Anne Scheidecker, Shanil Ebrahim, Vahid Ashoorion, Yasir Rehman, Patrick J. Hong, Stephanie Ross, Bradley C. Johnston, Regina Kunz, Xin Sun, Norman Buckley, Daniel I. Sessler, and Gordon H. Guyatt.** 2018. "Opioids for Chronic Noncancer Pain: A Systematic Review and Meta-analysis." *JAMA*, 320(23): 2448–2460.
- Callaway, Brantly, and Pedro H. C. Sant'Anna.** 2020. "Difference-in-Differences with Multiple Time Periods." *arXiv:1803.09015 [econ, math, stat]*.
- Case, Anne, and Angus Deaton.** 2015a. "Rising Morbidity and Mortality in Midlife among White Non-Hispanic Americans in the 21st Century." *Proceedings of the National Academy of Sciences*, 112(49): 15078–15083.
- Case, Anne, and Angus Deaton.** 2015b. "Suicide, Age, and Wellbeing: An Empirical Investigation." National Bureau of Economic Research Working Paper 21279.
- Case, Anne, and Angus Deaton.** 2017. "Mortality and Morbidity in the 21st Century." Brookings Papers on Economic Activity.
- Catan, Thomas, and Evan Perez.** 2012. "A Pain-Drug Champion Has Second Thoughts." Wall Street Journal.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer.** 2019. "The Effect of Minimum Wages on Low-Wage Jobs*." *The Quarterly Journal of Economics*, 134(3): 1405–1454.
- Center for Integrated Behavioral Health Policy.** 2011. "Methods Used by The Center for Integrated Behavioral Health Policy to Calculate Public Sector Costs of Substance Use Problems." George Washington University Medical Center.

- Centers for Disease Control.** 2018. "Calculating Total Daily Dose of Opioids For Safer Dosage."
- Centers for Disease Control.** 2019a. "CDC Advises Against Misapplication of the Guideline for Prescribing Opioids for Chronic Pain." <https://www.cdc.gov/media/releases/2019/s0424-advises-misapplication-guideline-prescribing-opioids.html>.
- Centers for Disease Control.** 2019b. "Prescribing Practices; Drug Overdose." <https://www.cdc.gov/drugoverdose/data/prescribing/prescribing-practices.html>.
- Centers for Disease Control and Prevention.** 2015. "CDC funding helps states combat prescription drug overdose epidemic." CDC Newsroom.
- Chandra, Amitabh, David Cutler, and Zirui Song.** 2012. "Who Ordered That? The Economics of Treatment Choices in Medical Care." *Handbook of Health Economics*, 2: 397–432.
- Chou, Roger, Jane C. Ballantyne, Gilbert J. Fanciullo, Perry G. Fine, and Christine Miaskowski.** 2009. "Research Gaps on Use of Opioids for Chronic Noncancer Pain: Findings From a Review of the Evidence for an American Pain Society and American Academy of Pain Medicine Clinical Practice Guideline." *The Journal of Pain*, 10(2): 147–159.e15.
- Chou, Roger, Judith A. Turner, Emily B. Devine, Ryan N. Hansen, Sean D. Sullivan, Ian Blazina, Tracy Dana, Christina Bougatsos, and Richard A. Deyo.** 2015. "The Effectiveness and Risks of Long-Term Opioid Therapy for Chronic Pain: A Systematic Review for a National Institutes of Health Pathways to Prevention Workshop." *Annals of Internal Medicine*, 162(4): 276–286.
- Christensen, Finn Bjarke.** 2004. "Lumbar Spinal Fusion. Outcome in Relation to Surgical Methods, Choice of Implant and Postoperative Rehabilitation." *Acta Orthopaedica Scandinavica. Supplementum*, 75(313): 2–43.
- de Chaisemartin, Clément, and Xavier D'Haultfœuille.** 2020. "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review*, 110(9): 2964–2996.
- Deiana, Claudio, and Ludovica Giua.** 2018. "The US Opidemic: Prescription Opioids, Labour Market Conditions and Crime." Working Paper.
- Delcher, Chris, Alexander C. Wagenaar, Bruce A. Goldberger, Robert L. Cook, and Mildred M. Maldonado-Molina.** 2015. "Abrupt Decline in Oxycodone-Caused Mortality after Implementation of Florida's Prescription Drug Monitoring Program." *Drug & Alcohol Dependence*, 150: 63–68.
- Dowell, Deborah, Tamara Haegerich, and Roger Chou.** 2019. "No Shortcuts to Safer Opioid Prescribing." *New England Journal of Medicine*, 380(24): 2285–2287.
- Eriksen, Jorgen, Per Sjogren, Eduardo Bruera, Ola Ekholm, and Niels K. Rasmussen.** 2006. "Critical Issues on Opioids in Chronic Non-Cancer Pain." *Pain*, 125(1): 172–179.
- Federation.** 1998. "Model Policy for the Use of Controlled Substances for the Treatment of Pain." Federation of State Medical Boards of the United States, Inc.
- Federation.** 2004. "Model Policy for the Use of Controlled Substances for the Treatment of Pain." Federation of State Medical Boards of the United States, Inc.
- Fink, David S., Julia P. Schleimer, Aaron Sarvet, Kiran K. Grover, Chris Delcher, Alvaro Castillo-Carniglia, June H. Kim, Ariadne E. Rivera-Aguirre, Stephen G. Henry, Silvia S. Martins, and**

- Magdalena Cerdá.** 2018. "Association Between Prescription Drug Monitoring Programs and Non-fatal and Fatal Drug Overdoses: A Systematic Review." *Annals of Internal Medicine*, 168(11): 783–790.
- Franco, Armando, Zachary Wagner, and Christopher Whaley.** 2019. "Doped up or Dropped Out: The Effects of Opioid Use on Workplace Absences." Working Paper.
- Franklin, Gary M., Bert D. Stover, Judith A. Turner, Deborah Fulton-Kehoe, Thomas M. Wickizer, and Disability Risk Identification Study Cohort.** 2008. "Early Opioid Prescription and Subsequent Disability among Workers with Back Injuries: The Disability Risk Identification Study Cohort." *Spine*, 33(2): 199–204.
- Frieden, Thomas R., and Debra Houry.** 2016. "Reducing the Risks of Relief – The CDC Opioid-Prescribing Guideline." *New England Journal of Medicine*, 374(16): 1501–1504.
- Friedman, Jerome, Trevor Hastie, and Robert Tibshirani.** 2009. *The Elements of Statistical Learning*. . Second ed., Springer, Berlin:Springer Series in Statistics.
- Furlan, Andrea D., Juan A. Sandoval, Angela Mailis-Gagnon, and Eldon Tunks.** 2006. "Opioids for Chronic Noncancer Pain: A Meta-Analysis of Effectiveness and Side Effects." *CMAJ: Canadian Medical Association Journal*, 174(11): 1589–1594.
- Garthwaite, Craig L.** 2012. "The Economic Benefits of Pharmaceutical Innovations: The Case of Cox-2 Inhibitors." *American Economic Journal: Applied Economics*, 4(3): 116–37.
- Gleason, Rae Marie, Kenneth L. Kirsh, Steven D. Passik, and Jan F. Chambers.** 2014. "Current Access to Opioids: Survey of Chronic Pain Patients." *Practical Pain Management*, 14(2).
- Goodman-Bacon, Andrew.** 2018. "Difference-in-Differences with Variation in Treatment Timing." National Bureau of Economic Research w25018, Cambridge, MA.
- Goodman-Bacon, Andrew.** 2019. "So You've Been Told to Do My Difference-in-Differences Thing: A Guide."
- Griffin, Beth Ann, Megan Schuler, Elizabeth Stuart, Stephen Patrick, Elizabeth McNeer, Rosanna Smart, David Powell, Bradley Stein, Terry Schell, and Rosalie Liccardo Pacula.** 2020. "Variation in Performance of Commonly Used Statistical Methods for Estimating Effectiveness of State-Level Opioid Policies on Opioid-Related Mortality." National Bureau of Economic Research w27029, Cambridge, MA.
- Gross, Douglas P., Brian Stephens, Yagesh Bhambhani, Mark Haykowsky, Geoff P. Bostick, and Saifudin Rashiq.** 2009. "Opioid Prescriptions in Canadian Workers' Compensation Claimants: Prescription Trends and Associations between Early Prescription and Future Recovery." *Spine*, 34(5): 525–531.
- Hackbarth, Glenn, Jon Christianson, and Mark Miller.** 2015. "Polypharmacy and Opioid Use among Medicare Part D Enrollees." In *Report to the Congress: Medicare and the Health Care Delivery System*. 117–138. Medicare Payment Advisory Commission.
- HCCI.** 2013. "Mental Health-Substance Use Services in Hospitals Up after Parity Law, Finds New Report." Health Care Cost Institute.
- Hoffman, Jan.** 2016. "Patients in Pain, and a Doctor Who Must Limit Drugs." *The New York Times*.

- Horwitz, Jill R., Corey Davis, Lynn McClelland, Rebecca Fordon, and Ellen Meara.** 2020. "The Importance of Data Source in Prescription Drug Monitoring Program Research." *Health Services Research*, 1475–6773.13548.
- Jones, Christopher M., Leonard J. Paulozzi, and Karin A. Mack.** 2014. "Sources of Prescription Opioid Pain Relievers by Frequency of Past-Year Nonmedical Use: United States, 2008-2011." *JAMA Internal Medicine*, 174(5): 802–803.
- Kaplovitch, Eric, Tara Gomes, Ximena Camacho, Irfan A. Dhalla, Muhammad M. Mamdani, and David N. Juurlink.** 2015. "Sex Differences in Dose Escalation and Overdose Death during Chronic Opioid Therapy: A Population-Based Cohort Study." *PLoS ONE*, 10(8): e0134550.
- Kasich, John.** 2015. "Empowering Prescribers & Pharmacists to Prevent Opiate Abuse."
- Kassed, Cheryl, Katharine Levit, and Megan Hambrick.** 2007. "Hospitalizations Related to Drug Abuse, 2005." Healthcare Cost and Utilization Project (HCUP).
- Kolodny, Andrew, David T. Courtwright, Catherine S. Hwang, Peter Kreiner, John L. Eadie, Thomas W. Clark, and G. Caleb Alexander.** 2015. "The Prescription Opioid and Heroin Crisis: A Public Health Approach to an Epidemic of Addiction." *Annual Review of Public Health*, 36(1): 559–574.
- Krebs, Erin E., Amy Gravely, Sean Nugent, Agnes C. Jensen, Beth DeRonne, Elizabeth S. Goldsmith, Kurt Kroenke, Matthew J. Bair, and Siamak Noorbaloochi.** 2018. "Effect of Opioid vs Nonopioid Medications on Pain-Related Function in Patients With Chronic Back Pain or Hip or Knee Osteoarthritis Pain: The SPACE Randomized Clinical Trial." *JAMA*, 319(9): 872–882.
- Krueger, Alan B.** 2017. "Where Have All the Workers Gone?: An Inquiry into the Decline of the U.S. Labor Force Participation Rate." *Brookings Papers on Economic Activity*, 2017(2): 1–87.
- Maclean, Johanna Catherine, Justine Mallatt, Christopher J. Ruhm, and Kosali Simon.** 2020. "Economic Studies on the Opioid Crisis: A Review." National Bureau of Economic Research Working Paper 28067.
- Madara, James L.** 2015. "AMA Letter to CDC."
- Madara, James L.** 2020. "AMA Urges CDC to Revise Opioid Prescribing Guideline."
- Mahmud, Mohammed A., Barbara S. Webster, Theodore K. Courtney, Simon Matz, James A. Tacci, and David C. Steffens.** 2000. "Clinical Management and the Duration of Disability for Work-Related Low Back Pain." *Journal of Occupational and Environmental Medicine / American College of Occupational and Environmental Medicine*, 42(12): 1178–1187.
- Manchikanti, Laxmaiah, Ricardo Vallejo, Kavita N. Manchikanti, Ramsin M. Benyamin, Sukdeb Datta, and Paul J. Christo.** 2011. "Effectiveness of Long-Term Opioid Therapy for Chronic Non-Cancer Pain." *Pain Physician*, 14(2): E133–156.
- McCance-Katz, Elinore, Andrew Kolodny, Michael R. Von Korff, and Jane C. Ballantyne.** 2012. "Why the FDA Should Revise Labels on Opioid Analgesics." *Medscape Neurology*.
- McCarberg, Bill.** 2016. "Re: Docket No. CDC-2015-0112; Proposed 2016 Guideline for Prescribing Opioids for Chronic Pain."
- McQuay, Henry.** 1999. "Opioids in Pain Management." *The Lancet*, 353(9171): 2229–2232.

- Meinhofer, Angelica.** 2015. "The War on Drugs: Estimating the Effect of Prescription Drug Supply-Side Interventions." Brown University Job Market Paper.
- Meinhofer, Angélica.** 2018. "Prescription Drug Monitoring Programs: The Role of Asymmetric Information on Drug Availability and Abuse." *American Journal of Health Economics*, 4(4): 504–526.
- Melzack, R.** 1990. "The Tragedy of Needless Pain." *Scientific American*, 262(2): 27–33.
- Miller, Joshua.** 2015. "Facing epidemic, Baker seeks to limit opioid prescriptions." Boston Globe.
- Morris, Brent J., Mitzi S. Laughlin, Hussein A. Elkousy, Gary M. Gartsman, and T. Bradley Edwards.** 2015. "Preoperative Opioid Use and Outcomes after Reverse Shoulder Arthroplasty." *Journal of Shoulder and Elbow Surgery*, 24(1): 11–16.
- Moyo, Patience, Linda Simoni-Wastila, Beth Ann Griffin, Eberechukwu Onukwugha, Donna Harrington, G. Caleb Alexander, and Francis Palumbo.** 2017. "Impact of Prescription Drug Monitoring Programs (PDMPs) on Opioid Utilization among Medicare Beneficiaries in 10 US States: Impact of PDMPs on Opioid-Prescribing in Medicare." *Addiction*, 112(10): 1784–1796.
- Nadeau, Stephen E., Jeffrey K. Wu, and Richard A. Lawhern.** 2021. "Opioids and Chronic Pain: An Analytic Review of the Clinical Evidence." *Frontiers in Pain Research*, 2: 44.
- Nerenberg, D. M., and J. Fudin.** 2010. "Maximum Daily Dose of Hydrocodone." *American Journal of Health-System Pharmacy*, 67(19): 1588–1589.
- Noble, Meredith, Jonathan R. Treadwell, Stephen J. Tregear, Vivian H. Coates, Philip J. Wiffen, Clarisse Akafomo, Karen M. Schoelles, and Roger Chou.** 2010. "Long-Term Opioid Management for Chronic Noncancer Pain." *Cochrane Database of Systematic Reviews*, 20(1).
- Pardo, Bryce.** 2017. "Do More Robust Prescription Drug Monitoring Programs Reduce Prescription Opioid Overdose?: Prescription Drug Monitoring and Opioid Overdoses." *Addiction*, 112(10): 1773–1783.
- Park, Sujeong, and David Powell.** 2020. "Is the Rise in Illicit Opioids Affecting Labor Supply and Disability Claiming Rates?" National Bureau of Economic Research w27804, Cambridge, MA.
- Patrick, Stephen W., Carrie E. Fry, Timothy F. Jones, and Melinda B. Buntin.** 2016. "Implementation Of Prescription Drug Monitoring Programs Associated With Reductions In Opioid-Related Death Rates." *Health Affairs*, 35(7): 1324–1332.
- Paulozzi, Leonard J., Edwin M. Kilbourne, and Hema A. Desai.** 2011. "Prescription Drug Monitoring Programs and Death Rates from Drug Overdose." *Pain Medicine*, 12(5): 747–754.
- PDMP Center of Excellence.** 2014. "Briefing on PDMP Effectiveness." Brandeis University.
- Portenoy, Russell K., and Kathleen M. Foley.** 1986. "Chronic Use of Opioid Analgesics in Non-Malignant Pain: Report of 38 Cases." *Pain*, 25(2): 171–186.
- Powell, David, and Seth A. Seabury.** 2014. "Medical Care Spending and Labor Market Outcomes: Evidence from Workers' Compensation Reforms." Social Science Research Network SSRN Scholarly Paper ID 2401860, Rochester, NY.
- Reifler, Liza M., Danna Droz, J. Elise Bailey, Sidney H. Schnoll, Reginald Fant, Richard C. Dart, and Becki Bucher Bartelson.** 2012. "Do Prescription Monitoring Programs Impact State Trends in Opioid Abuse/Misuse?" *Pain Medicine*, 13(3): 434–442.

- Reisman, Richard M., Preen J. Shenoy, Adam J. Atherly, and Christopher R. Flowers.** 2009. "Prescription Opioid Usage and Abuse Relationships: An Evaluation of State Prescription Drug Monitoring Program Efficacy." *Substance Abuse: Research and Treatment*, 3: SART.S2345.
- Rosenblum, Andrew, Lisa A. Marsch, Herman Joseph, and Russell K. Portenoy.** 2008. "Opioids and the Treatment of Chronic Pain: Controversies, Current Status, and Future Directions." *Experimental and Clinical Psychopharmacology*, 16(5): 405–416.
- Rosenblum, Keith E.** 2012. "Opioids Wreak Havoc on Workers' Compensation Costs." Lockton Companies.
- Rutkow, Lainie, Hsien-Yen Chang, Matthew Daubresse, Daniel W. Webster, Elizabeth A. Stuart, and G. Caleb Alexander.** 2015. "Effect of Florida's Prescription Drug Monitoring Program and Pill Mill Laws on Opioid Prescribing and Use." *JAMA Internal Medicine*, 175(10): 1642–1649.
- SAMHSA.** 2012. "Results from the 2011 National Survey on Drug Use and Health: Summary of National Findings." Substance Abuse and Mental Health Services Administration HHS Publication No. (SMA) 12-4713 NSDUH Series H-44, Rockville, MD.
- Savych, Bogdan, David Neumark, and Randall Lea.** 2019. "Do Opioids Help Injured Workers Recover and Get Back to Work? The Impact of Opioid Prescriptions on Duration of Temporary Disability." *Industrial Relations: A Journal of Economy and Society*, 58(4): 549–590.
- Sjogren, Per, Morten Gronbaek, Vera Peuckmann, and Ola Ekholm.** 2010. "A Population-Based Cohort Study on Chronic Pain: The Role of Opioids." *The Clinical Journal of Pain*, 26(9): 763–769.
- Slavova, Svetla, Daniella Bradley O'Brien, Kathleen Creppage, Dan Dao, Anna Fondario, Elizabeth Haile, Beth Hume, Thomas W. Largo, Claire Nguyen, Jennifer C. Sabel, and Dagan Wright.** 2015. "Drug Overdose Deaths: Let's Get Specific." *Public Health Reports*, 130(4): 339–342.
- Smith, Peter Andrey.** 2013. "Big Pharma's Private War on Drugs." The American Prospect.
- Sun, Liyang, and Sarah Abraham.** 2020. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." *arXiv:1804.05785 [econ]*.
- Thomas, D., J. Frascella, T. Hall, W. Smith, W. Compton, W. Koroshetz, J. Briggs, P. Grady, M. Somerman, and N. Volkow.** 2015. "Reflections on the Role of Opioids in the Treatment of Chronic Pain: A Shared Solution for Prescription Opioid Abuse and Pain." *Journal of Internal Medicine*, 278(1): 92–94.
- Tian, Terrence Y., Ianita Zlateva, and Daren R. Anderson.** 2013. "Using electronic health records data to identify patients with chronic pain in a primary care setting." *Journal of the American Medical Informatics Association: JAMIA*, 20(e2): e275–280.
- Twillman, R.** 2006. "Impact of Prescription Monitoring Programs on Prescription Patterns and Indicators of Opioid Abuse." *The Journal of Pain*, 7(4): S6–S6.
- US Department of Health and Human Services.** 2015. "HHS Takes Strong Steps to Address Opioid-Drug Related Overdose, Death and Dependence." HHS Press Office.
- U.S. Department of Justice.** 2020. "Justice Department Announces Global Resolution of Criminal and Civil Investigations with Opioid Manufacturer Purdue Pharma and Civil Settlement with Members of the Sackler Family." <https://www.justice.gov/opa/pr/justice-department-announces-global-resolution-criminal-and-civil-investigations-opioid>.

- Vogt, Molly T., C. Kent Kwok, Doris K. Cope, Thaddeus A. Osial, Michael Culyba, and Terence W. Starz.** 2005. "Analgesic Usage for Low Back Pain: Impact on Health Care Costs and Service Use." *Spine*, 30(9): 1075–1081.
- Volinn, Ernest, Jamison D. Fargo, and Perry G. Fine.** 2009. "Opioid Therapy for Nonspecific Low Back Pain and the Outcome of Chronic Work Loss." *Pain*, 142(3): 194–201.
- Vowles, Kevin E., Mindy L. McEntee, Peter Siyahhan Julnes, Tessa Frohe, John P. Ney, and David N. van der Goes.** 2015. "Rates of Opioid Misuse, Abuse, and Addiction in Chronic Pain: A Systematic Review and Data Synthesis." *PAIN*, 156(4): 569–576.
- Webster, Barbara S., Santosh K. Verma, and Robert J. Gatchel.** 2007. "Relationship between Early Opioid Prescribing for Acute Occupational Low Back Pain and Disability Duration, Medical Costs, Subsequent Surgery and Late Opioid Use." *Spine*, 32(19): 2127–2132.
- White House.** 2005. "National Drug Control Strategy."
- Zusman, Edie E.** 2012. "HCAHPS Replaces Press Ganey Survey as Quality Measure for Patient Hospital Experience:." *Neurosurgery*, 71(2): N21–N24.

Tables

Table 1: Summary Statistics: ARCOS

MME per Resident per Quarter	Mean	Standard Deviation	Minimum	Median	Maximum
Oxycodone	58.8	31.3	12.4	53.5	284.8
<i>Other Schedule II :</i>					
Fentanyl	26.9	8.2	8.1	26.8	52.5
Hydromorphone	3.7	2.1	0.4	3.2	16.3
Meperidine	1.2	0.9	0.1	0.9	5.3
Morphine	16.1	6.9	4.4	15.0	41.4
Total Schedule II	106.7	42.7	31.1	101.7	343.8
Hydrocodone	23.2	14.1	3.3	20.5	84.8
<i>Other Schedule III-V :</i>					
Codeine	1.4	0.4	0.6	1.3	4.5
Total Schedule III-V	24.6	14.0	4.1	22.0	86.3

Notes: Source is Drug Enforcement Administration Automated Reports and Consolidated Orders System (ARCOS). Summary statistics for 38 states in sample depicted in Table A1. Data spans years 2001 to 2013. Quantities converted to morphine equivalents using a conversion factor, morphine milligram equivalents, or MME (Centers for Disease Control, 2018). Sample: 1,976 state-quarters. Oxycodone is classified as Schedule II, while hydrocodone was classified as Schedule III during this period. Schedule II and III drugs are both defined as having “a currently accepted medical use in treatment in the United States,” and a “potential for abuse and dependence.” Schedule III drugs are distinguished from Schedule II as less abusable, and less likely to produce physical and psychological dependence. Schedule III drugs are easier to prescribe and obtain.

Table 2: Summary Statistics: Marketscan*Panel A: Percentage of enrollees by year with prescription for opioid pain relievers*

Year	Oxycodone	Hydrocodone	Schedule II excluding oxycodone	Schedule III & below excluding hydrocodone	All opioids	Total enrollees
2005	6.0	15.9	1.0	9.6	25.3	822,167
2006	5.8	15.4	0.9	8.6	24.2	1,057,487
2007	5.9	15.7	0.9	8.6	24.5	1,221,408
2008	6.2	15.7	0.8	9.7	25.4	1,579,009
2009	6.5	15.6	0.8	9.7	25.6	2,079,528
2010	6.4	14.9	0.8	8.7	24.2	2,527,267
2011	6.5	15.3	0.9	8.1	24.3	2,686,926
2012	6.2	14.7	0.8	7.8	23.5	2,774,767
All Years	14.0	29.4	2.0	18.4	41.9	4,944,277

Panel B: Dosage by time elapsed in opioid therapy episode

Days elapsed	Hydrocodone (average MME per day in script)	Oxycodone (average MME per day in script)
Days 0 to 15	33.3	39.6
Days 15 to 30	27.8	40.2
Days 30 to 60	28.4	45.6
Days 60 to 90	31.1	56.3
Days 90 to 120	32.8	65.1
Days 120 to 150	33.7	69.3
Days 150 to 180	34.3	73.3
Days 180 to 365	36.0	83.0
Days 365 and after	37.9	103.4

Notes: Panel A displays the percent of enrollees in 38-state sample with an oxycodone, hydrocodone, etc. prescription in each year, as well as over all years spent in the Marketscan prescription drug claims dataset. Panel B displays the average MME prescribed for prescriptions that are earlier or later in an opioid therapy episode, where an opioid therapy episode is as defined in the text.

Table 3: The Effect of PDMP Introduction on all Opioids Distributed (ARCOS)

	(1)	(2)	(3)	(4)
	Total Schedule II	Oxycodone	Total Schedule II excl. Oxycodone	Hydrocodone
1(PDMP) _{-st}	-11.15** (5.036)	-10.68** (4.858)	-0.471 (1.102)	2.466* (1.412)
Controls:				
State Controls	Y	Y	Y	Y
State FE	Y	Y	Y	Y
Quarter FE	Y	Y	Y	Y
Clusters	38	38	38	38
Observations	1,976	1,976	1,976	1,976
Mean	106.7	58.8	47.9	23.2

Notes: The table reports estimates of the impact of the introduction of a PDMP on all opioids distributed, as monitored by the ARCOS system, estimated according to Equation 2. The dependent variables are the sum of Morphine Milligram Equivalent (MME) per resident per quarter, for four categories of opioids. Column (1) represents the impact on all Schedule II drugs: oxycodone, fentanyl, hydromorphone, meperidine, and morphine. Column (2) is for oxycodone only; column (3) is all other Schedule II drugs, and column (4) is hydrocodone, the highest-volume opioid below Schedule II. State controls include state-level population, unemployment rate, and percentage of population over 60. Robust standard errors, clustered at the state level, in parentheses.

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

Table 4: The Effect of PDMP Introduction on Overdose Deaths (CDC Vital Statistics)

	(1)	(2)
	Opioid Analgesic Overdose Deaths (T40.2)	All Opioid Overdose Deaths (T40.1-T40.4,T40.6)
$1(\text{PDMP})_{st}$	-0.108** (0.0439)	-0.0816 (0.0536)
Controls:		
County Controls	Y	Y
County FE	Y	Y
Quarter FE	Y	Y
Clusters	38	38
Observations	117,856	117,856
Counties	2,267	2,267
Mean	0.62	1.37

Notes: The table reports estimates of the impact of the introduction of a PDMP on opioid overdose deaths. Each column reports the dynamic ATT for the impact of PDMP introduction, as described in Section 5.2 and Appendix C. Dependent variables are $\ln(\text{deaths} + 1)$ in a county for a given quarter. Drug overdoses are identified using underlying cause of death codes X40-X44, X60-X64, X85, and Y10-Y14. “Opioid Analgesic Overdoes Deaths” are defined as deaths that list multiple cause of Death code T40.2. “All Opioid Overdoes Deaths” are defined as deaths with multiple cause of death codes T40.1, T40.2, T40.3, T40.4 and/or T40.6; categories are described above in Section 3.3. The baseline specification includes county-level controls for the natural log of population, unemployment rate, and percentage of population over 60, and observations are weighted according to county population. Means are unweighted and listed in levels (count of deaths in a county-quarter). Robust standard errors, clustered at the state level, in parentheses.

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

Table 5: The Effect of PDMP Introduction on Episodes of Abuse, for all enrollees and chronic pain patients (Marketscan)

	(1)	(2)	(3)
1(PDMP) <i>_st</i>	-0.0606** (0.0273)	-0.0314 (0.0304)	-0.0842** (0.0415)
Sample	Full	Chronic pain	Full excl. chronic pain
Model	Logistic	Logistic	Logistic
Individual Controls	Y	Y	Y
Individual FE	N	N	N
Entry-cohort FE	Y	Y	Y
Quarter FE	Y	Y	Y
Observations	53,364,705	5,824,330	47,194,901
Clusters	38	38	38
Individuals	4,933,017	337,955	4,573,291
Mean	0.00156	0.00763	0.000819

Notes: Models are estimated according to Equation 6. The dependent variable is a binary outcome variable indicating an episode of abuse identified in Marketscan, and construction is described in Section 5.3. The chronic pain patient sample is selected as described in Section 3.2.1 and Appendix E. Means are in levels and reveal the baseline opioid use disorder rate in the full, chronic pain, and non-chronic pain sample. Robust standard errors, clustered at the state level, in parentheses.

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

Table 6: The Effect of PDMP Introduction on Schedule II Opioid Prescriptions Filled (Marketscan)

	Quantity			Probability	
	(1)	(2)	(3)	(4)	(5)
1(PDMP) <i>_st</i>	-0.00482*** (0.00160)	-0.0898** (0.0432)	-6.431** (2.708)	-0.000562** (0.000226)	-0.0226** (0.00978)
Model	Log-linear	PPML	Linear	LPM	Logistic
Observations	53,520,024	53,520,024	53,520,024	53,520,024	53,520,024
Clusters	38	38	38	38	38
Individuals	4,944,277	4,944,277	4,944,277	4,944,277	4,944,277
Individual Controls	Y	Y	Y	Y	Y
Individual FE	Y	N	Y	Y	N
Entry-cohort FE	N	Y	N	N	Y
Quarter FE	Y	Y	Y	Y	Y
Mean	0.155	60.99	60.99	0.0247	0.0247

Notes: Models are estimated according to Equation 8. “Schedule II opioid prescriptions” in the Marketscan sample are filled prescription claims for oxycodone, fentanyl, hydromorphone, meperidine, and morphine. For Model (1), the dependent variable is a concave transformation of Morphine Milligram Equivalent (MME) for Schedule II drugs per quarter for each individual enrollee, $\ln(MME + 1)$. Model (2) is a Poisson pseudo-maximum likelihood regression for MME of Schedule II opioids prescribed for each individual enrollee in a given quarter, implemented using the `ppmlhdf` package in Stata, and Model (3) is an OLS regression for MME of opioids prescribed (not transformed). Models (4) and (5) are probability models capturing the extensive margin of prescribing, using a binary outcome variable for whether an individual was prescribed a Schedule II opioid in a given quarter. Each specification includes individual or entry-cohort fixed effects and quarter fixed effects, and individual-level controls (age, employment type/status [salaried, union, full time, etc.], and industry [oil and gas, manufacturing, etc.]). Robust standard errors, clustered at the state level, in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: The Effect of PDMP Introduction on the Distribution of Opioid Prescribing (Marketscan)

	(1)	(2)	(3)
Panel A: New Schedule II opioid prescription filled in quarter			
(A-1) Initial Schedule II prescription	-0.000262** (0.000117)	-0.0191** (0.00914)	0.0165
Panel B: Enrollee dispensed amount between thresholds in quarter			
(B-1) No Schedule II opioids	0.000562** (0.000226)	0.0226** (0.00978)	0.975
(B-2) > 0 to 224 MME	-0.0000453 (0.0000537)	-0.0160 (0.0122)	0.00594
(B-3) 225 to 1,799 MME	-0.000115 (0.000153)	-0.00769 (0.0113)	0.0138
(B-4) 1,800 to 4,499 MME	-0.000104 (0.0000632)	-0.0281 (0.0192)	0.00204
(B-5) 4,500 to 8,099 MME	-0.0000779** (0.0000368)	-0.0616** (0.0279)	0.000996
(B-6) >8,100 MME	-0.000221*** (0.0000767)	-0.109** (0.0426)	0.00189
Model	LPM	Logistic	Means
Individual Controls	Y	Y	
Individual FE	Y	N	
Entry-cohort FE	N	Y	
Quarter FE	Y	Y	

Notes: Each model estimates, according to Equation 8, the probability that in a given quarter, an enrollee in the Marketscan sample falls into one of the above categories after the introduction of a PDMP, reporting the point estimate on $1(\text{PDMP})_{st}$. In Panel A, the dependent variable, “initial Schedule II prescription,” is a binary variable indicating that an enrollee has a prescription that was not preceded by another Schedule II prescription in the two previous quarters. In Panel B, the dependent variable in each panel is a binary indicator for whether an enrollee was prescribed an amount between the listed thresholds in a quarter. Thresholds are set based on the CDC prescribing guidelines discussed in Footnote 2. Panel B-2 represents dosing that is likely to be an acute prescription. Panel B-3 includes dosing above the acute threshold, and up to a maximum of 20 MME/day for the quarter, or a low chronic dose. Panel B-4 includes dosing between 20 and 50 MME per day, where 50 MME per day is the CDC’s recommended maximum dose for most chronic prescriptions. Panel B-5 includes dosing between the 50 MME per day guideline and the 90 MME per day recommended maximum for all chronic prescriptions. Panel B-6 represents dosing exceeding the CDC’s recommended 90 MME daily maximum. Column (3) reports the mean of each binary dependent variable. For Columns (1) and (2), Panel (A), the sample comprises 4,359,108 individuals, 38 states, and 44,909,146 observations; for Column (1), Panel (B), the sample is the same as in Table 6. For Column (2), Panel (B), the sample varies slightly in each specification from 53,473,830 to 53,520,024 observations, as a few small entry-cohorts do not have any positive observations. Robust standard errors, clustered at the state level, in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: The Effect of PDMP Introduction on Opioid Prescribing, Inpatient and Outpatient Spending, and Work Output for all enrollees and chronic pain patients (Marketscan)

	All Enrollees			Chronic Pain Sample		
	(1) ln(MME +1)	(2) ln(spending +1)	(3) days absent	(4) ln(MME +1)	(5) ln(spending +1)	(6) days absent
1(PDMP) _{st}	-0.00482*** (0.00160)	0.0197 (0.0179)	0.0512* (0.0306)	-0.0145** (0.00712)	0.0609** (0.0232)	0.0622** (0.0313)
Model	Log-linear	Log-linear	PPML	Log-linear	Log-linear	PPML
Individual Controls	Y	Y	Y	Y	Y	Y
Individual FE	Y	Y	N	Y	Y	N
Entry-cohort FE	N	N	Y	N	N	Y
Quarter FE	Y	Y	Y	Y	Y	Y
Observations	53,520,024	53,520,057	53,520,035	5,870,917	5,870,917	5,870,916
Clusters	38	38	38	38	38	38
Individuals	4,944,277	4,944,277	4,944,277	340,660	340,660	340,660
Mean	60.99	\$899.97	0.732	318.0	\$2,987.82	3.1

Notes: Models are estimated according to Equation 8. The chronic pain sample is constructed as described in Section 3.2.1 and Appendix E. The dependent variables are ln(MME schedule II opioids + 1), ln(inpatient and outpatient spending + 1), and count of days absent under workers compensation and short term disability, per quarter for each enrollee. Robust standard errors, clustered at the state level, in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 9: The Effect of PDMP Introduction on Total Costs (Inpatient, Outpatient, Drug, Lost Wages, WC Medical) for all enrollees and chronic pain patients (Marketscan)

	(1)	(2)
1(PDMP) _{st}	0.0133 (0.0145)	0.0451** (0.0195)
Sample	Full	Chronic Pain
Individual Controls	Y	Y
Individual FE	Y	Y
Quarter FE	Y	Y
Observations	53,518,181	5,870,689
Clusters	38	38
Individuals	4,944,255	340,660
Mean	\$1,215.70	\$3,919.53
Implied increased costs (billions)	\$10.9	\$8.2

Notes: Models are estimated according to Equation 8. The chronic pain sample is constructed as described in Section 3.2.1 and Appendix E. Dependent variable is ln(total inpatient spending, outpatient spending, drug spending, imputed lost wages under workers' compensation and short-term disability, and imputed medical costs accrued to workers' compensation insurer +1). Means are listed in levels (\$). Implied total costs in billions are calculated as described in Section 7.2. Robust standard errors, clustered at the state level, in parentheses.

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

Table 10: The Effect of PDMP Introduction on Opioid Prescribing, by Demographic Characteristics (Marketscan)

	(1)	(2)	(3)
Pain conditions			
Chronic pain sample	-0.0197* (0.0108)	340,660	0.0817
Full sample excl. chronic pain	-0.0222** (0.0101)	4,603,617	0.0176
Gender			
Male	-0.0327*** (0.00847)	2,831,565	0.0213
Female	-0.00825 (0.0138)	2,112,747	0.0297
Age			
17-28	-0.00951 (0.0195)	1,133,700	0.0172
29-40	-0.0121 (0.0102)	1,849,690	0.0231
41-52	-0.0309*** (0.0113)	1,854,040	0.0264
53-64	-0.0266** (0.0118)	1,096,376	0.0290
Employment			
Union/Hourly/Manufacturing	-0.0346*** (0.0130)	3,189,149	0.0257
All others	0.00570 (0.0110)	2,054,485	0.0229
Urbanicity			
Urban (in MSA)	-0.0250** (0.0101)	4,509,599	0.0247
Rural	-0.0252 (0.0191)	480,457	0.0241
Model	Logistic	Individuals	Means
Individual Controls	Y		
Individual FE	N		
Entry-cohort FE	Y		
Quarter FE	Y		

Notes: Each model estimates, according to Equation 8, the probability that in a given quarter, an enrollee in the Marketscan sample has an opioid prescription after the introduction of a PDMP, reporting the point estimate on $1(\text{PDMP})_{st}$. Each specification is a logistic regression run on a split sample selected based on the identified characteristic. The chronic pain sample is constructed as described in Section 3.2.1 and Appendix E; all other samples are constructed based on the provided Marketscan demographic indicators. Robust standard errors, clustered at the state level, in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

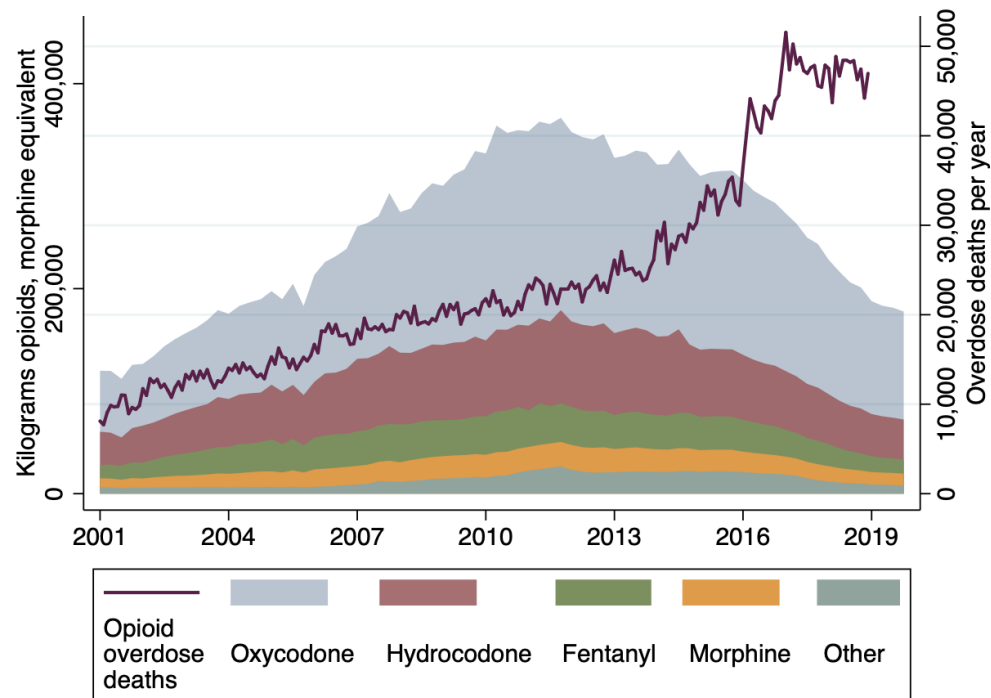
Table 11: The Effect of PDMP Introduction on Opioid Overdose Deaths, by Demographic Characteristics (CDC Vital Statistics)

	(1)	(2)
Gender		
Male	-0.0633* (0.0332)	8.03
Female	-0.139* (0.0749)	5.15
Age		
17-28	-0.0611 (0.0537)	5.89
29-40	-0.110 (0.0748)	9.91
41-52	-0.0842 (0.0589)	14.03
53-64	-0.0487 (0.0410)	8.87
Education		
No high school diploma	-0.107 (0.0680)	12.75
High school but no 4-year college degree	-0.140*** (0.0066)	11.64
College or post-graduate degree	-0.0148 (0.0502)	3.49
Urbanicity		
Urban (in MSA)	-0.126** (0.0570)	6.41
Rural	-0.0122 (0.0230)	7.27
Model	Log-linear	Deaths per pop per quarter
County Controls	Y	
County FE	Y	
Quarter FE	Y	

Notes: Each model reports the dynamic ATT for the impact of PDMP introduction, as described in Section 5.2 and Appendix C. Dependent variables are $\ln(\text{deaths with T40.2} + 1)$ in a county for a given quarter. The baseline specification includes county-level controls for the natural log of population, unemployment rate, and percentage of population over 60, and observations are weighted according to county population. Column (2) reports the opioid overdose death rate per capita for each demographic (measured from 2001-2013 in the 38-state sample). Robust standard errors, clustered at the state level, in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

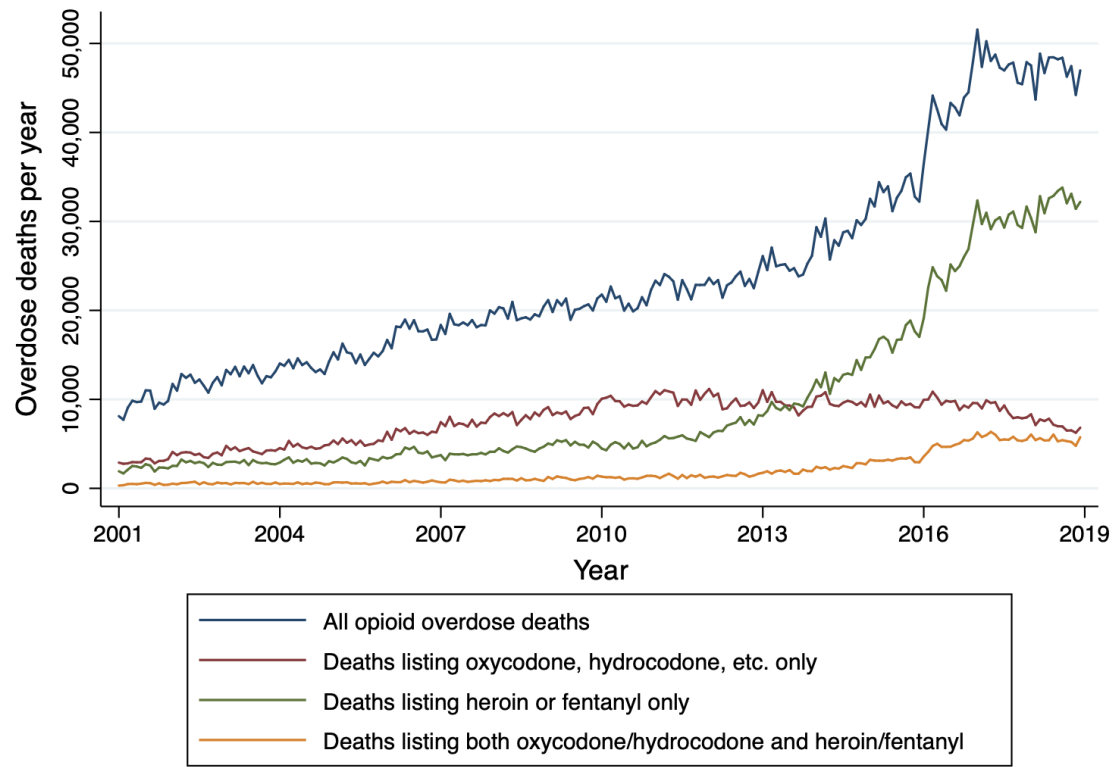
Figures

Figure 1: Opioids dispensed via pharmacies, hospitals, practitioners, etc., overlaid with opioid overdose deaths, from 2001-2019




Notes: Source: Data extracted from Drug Enforcement Administration Automated Reports and Consolidated Ordering System (ARCOS) reports. Figure depicts total kilograms of opioids dispensed in morphine equivalents, and the proportion of that total accounted for by oxycodone, hydrocodone, etc. Each opioid is converted to morphine equivalents using a conversion factor, morphine milligram equivalents, or MME ([Centers for Disease Control, 2018](#)). Opioid overdose deaths are from CDC National Center for Health Statistics Vital Statistics mortality records. Opioid overdose deaths are deaths with underlying cause of death code of X40-X44, X60-X64, X85, or Y10-Y14 (indicating a drug overdose), and multiple causes of death codes of T40.0-T40.4 or T40.6 (which indicate opioids or heroin, specifically, were listed on the death certificate).

Figure 2: Opioid Overdose Deaths, 2001-2018 (CDC Vital Statistics)




Notes: Source: CDC National Center for Health Statistics Vital Statistics records. This figure depicts the changing composition of opioid overdose deaths over time. “Deaths listing oxycodone, hydrocodone, etc. only” are deaths only containing multiple cause of death code T40.2, which includes oxycodone, hydrocodone, morphine, etc. “Deaths listing heroin or fentanyl only” contains any death listing T40.1 (heroin) or T40.4 (fentanyl or other synthetic opioids), or both, but not other opioids. “Deaths listing both oxycodone/hydrocodone and heroin/fentanyl” are deaths listing both T40.2 and either T40.1 or T40.4. While opioid overdose deaths related only to prescription opioid pain relievers long drove total opioid overdose deaths, in recent years heroin, and then later, fentanyl, became the main driver.

Figure 3: The Effect of PDMP Introduction on all Schedule II Opioids Distributed (ARCOS)




Notes: ARCOS event study analogue of baseline specification for Schedule II opioids displayed in Table 3, Model (1). The model is estimated according to Equation 1. Date of implementation is date PDMP began prescription record collection, as described in Section 4.1. Robust standard errors clustered at state level. Thin bars represent 95% confidence intervals and thick bars represent 90% confidence intervals.

Figure 4: The Effect of PDMP Introduction on Opioid Overdose Deaths (CDC Vital Statistics)




Notes: CDC event study, estimated according to Equation 3, for opioid overdose deaths linked to semi-synthetic opioids such as oxycodone, hydrocodone, and morphine (identified by multiple causes of death code T40.2). Estimated using an OLS log-linear model where the dependent variable is $\ln(\text{opioid [T40.2] overdose deaths in county} + 1)$. The baseline specification includes controls for the natural log of population, unemployment rate, and percentage of population over 60, and observations are weighted according to county population. The results of this specification are reproduced in Table B3. Robust standard errors clustered at state level. Thin bars represent 95% confidence intervals and thick bars represent 90% confidence intervals.

Figure 5: The Effect of PDMP Introduction on Episodes of Opioid Use Disorder (Marketscan)



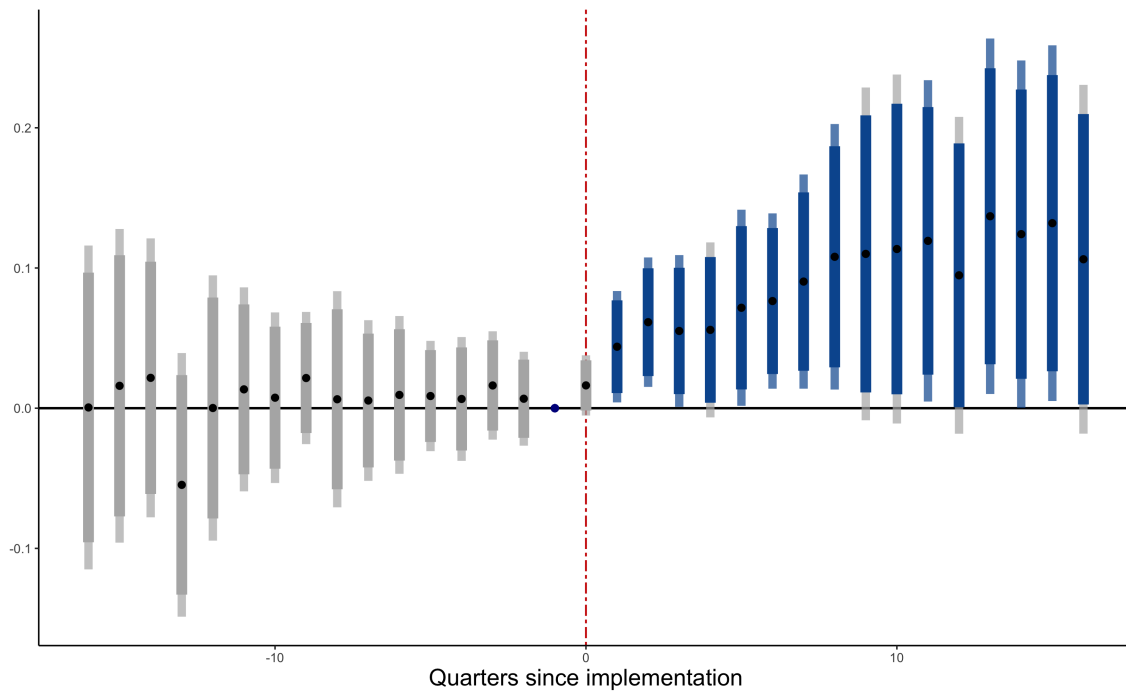
Notes: Marketscan event study analogue for Model (1) in Table 5, estimated according to Equation 5. Includes the full sample of patients. Robust standard errors clustered at state level. Thin bars represent 95% confidence intervals and thick bars represent 90% confidence intervals.

Figure 6: The Effect of PDMP Introduction on Schedule II Opioids Dispensed (Marketscan)




Notes: Marketscan event study analogue for Model (1) in Table 6, estimated according to Equation 7. Includes the full sample of patients. Robust standard errors clustered at state level. Thin bars represent 95% confidence intervals and thick bars represent 90% confidence intervals.

Figure 7: The Effect of PDMP Introduction on Total Inpatient and Outpatient Spending (Marketscan)




Notes: Marketscan event study analogue for Model (5) in Table 8, estimated according to Equation 7. The chronic pain sample is constructed as described in Section 3.2.1 and Appendix E. Robust standard errors clustered at state level. Thin bars represent 95% confidence intervals and thick bars represent 90% confidence intervals.

Figure 8: The Effect of PDMP Introduction on Days Absent from Work (Marketscan)



Notes: Marketscan event study analogue for Model (6) in Table 8, estimated according to Equation 7. The chronic pain sample is constructed as described in Section 3.2.1 and Appendix E. Robust standard errors clustered at state level. Thin bars represent 95% confidence intervals and thick bars represent 90% confidence intervals.

Figure 9: The Effect of PDMP Introduction on Total Costs (Marketscan)



Notes: Marketscan event study analogue for Model (2) in Table 9, estimated according to Equation 7. The chronic pain sample is constructed as described in Section 3.2.1 and Appendix E. Robust standard errors clustered at state level. Thin bars represent 95% confidence intervals and thick bars represent 90% confidence intervals.

Appendix A Prescription Drug Monitoring Programs

Table A1: Dates of Implementation of State PDMP (Main Sample)


Wyoming	Jan-2004	California	Jan-2009	South Dakota	Mar-2012
Maine	Jan-2005	Louisiana	Jan-2009	New Mexico	Aug-2012
Mississippi	Dec-2005	Iowa	Mar-2009	Delaware	Aug-2012
Virginia	Jun-2006	Vermont	Apr-2009	Montana	Oct-2012
Oklahoma	Jul-2006	Minnesota	Apr-2010	Arkansas	Mar-2013
Indiana	Jan-2007	Massachusetts	Aug-2010	Wisconsin	May-2013
North Dakota	Jan-2007	Kansas	Apr-2011	Georgia	Jul-2013
Alabama	Apr-2007	Oregon	Sep-2011	Maryland	Jan-2014
North Carolina	Oct-2007	Florida	Oct-2011	Nebraska	
Colorado	Feb-2008	Ohio	Oct-2011	Pennsylvania	
South Carolina	Jun-2008	Alaska	Jan-2012	New Hampshire	
Connecticut	Jul-2008	New Jersey	Jan-2012	Missouri	
Arizona	Dec-2008	Washington	Jan-2012		

Table A2: Year of Implementation of Early State Controlled Substances Monitoring

New York	1973
Texas	1989
Tennessee	1990
West Virginia	1995
Hawaii	1996
Utah	1997
Nevada	1997
Idaho	1998
Michigan	1998
Kentucky	1999
Illinois	1999
Rhode Island	2001

Notes: Source: National Alliance for Model State Drug Laws (NAMSDL).

Figure A2: Grams Morphine Equivalent Distributed per Person in 38 Sample States versus States with Early Controlled Substances Efforts



Notes: Source: DEA ARCOS.

Appendix B Robustness of key results to alternate specifications

Table B1: The Effect of PDMP Introduction on Schedule II Opioids Distributed (ARCOS): Robustness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
1(PDMP) _{st}	-11.15** (5.036)	-10.59** (5.096)	-12.10** (4.682)	-21.01* (10.56)	-8.053** (3.252)	-9.372* (5.259)	-10.78* (5.447)	-12.09** (5.359)	-7.670* (4.206)	-9.315 (5.606)	-8.131** (3.350)
Observations	1,976	1,976	1,976	1,976	1,924	1,976	1,872	1,872	2,340	1,716	1,976
Clusters	38	38	38	38	37	38	36	36	45	33	38
Mean	106.7	106.7	106.7	107.3	102.1	106.7	106.4	106.4	104.1	106.7	106.7
State controls	Y	N	Y	Y	Y	Y	Y	Y	Y	Y	Y
State FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Quarter FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
State linear trends	N	N	Y	N	N	N	N	N	N	N	N
Population-weighted	N	N	N	Y	Y	N	N	N	N	N	N
Excl. Florida	N	N	N	N	Y	N	N	N	N	N	N
OH and NM recode	N	N	N	N	N	Y	N	N	N	N	N
Drop OH and NM	N	N	N	N	N	N	Y	N	N	N	N
Implement after 2005	N	N	N	N	N	N	N	Y	N	N	N
Implement after 1997	N	N	N	N	N	N	N	N	Y	N	N
Horwitz legal coding	N	N	N	N	N	N	N	N	N	Y	N
DCDH'20 heterogeneity	N	N	N	N	N	N	N	N	N	N	Y

Notes: This table displays robustness checks corresponding to the baseline ARCOS specification for Schedule II opioids found in Table 3. State controls include population, unemployment rate, and percentage of population over 60. In Column (6), Ohio and New Mexico are coded with the first NAMSDL date of implementation in the study period rather than the second. In column (10), I use the legal coding from Horwitz et al. (2020) (date “Modern system operational”), while dropping all states that they indicate had an early PDMP that was legislated and enacted prior to 2003. In column (11) I report the average effect of 16 quarters of dynamic treatment effects estimated using the de Chaisemartin and D’Haultfoeuille (2020) `did_multiplegt` package, which computes unbiased and consistent estimators and is robust to heterogeneous treatment effects. Robust standard errors, clustered at the state level, in parentheses.

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

Table B2: Opioid overdose deaths (CDC Vital Statistics): Main results and robustness checks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Two-way fixed effects model							
1(PDMP) _{st}	-0.0713 (0.0453)	-0.0446 (0.0436)	-0.0709 (0.0474)	-0.0805 (0.0485)	-0.0718 (0.0459)	-0.0734* (0.0383)	-0.0645 (0.0610)
Panel B: Dynamic ATT							
	-0.108** (0.0439)	-0.0956* (0.0519)	-0.0836** (0.04037)	-0.0978** (0.0483)	-0.109** (0.0404)	-0.0906** (0.0429)	-0.0671* (0.0358)
Observations	117,856	114,372	117,856	111,564	115,828	137,200	101,840
Clusters	38	37	38	36	36	45	33
Mean	0.62	0.56	0.62	0.60	0.62	0.61	0.51
Log-linear model	Y	Y	Y	Y	Y	Y	Y
County controls	Y	Y	Y	Y	Y	Y	Y
County FE	Y	Y	Y	Y	Y	Y	Y
Quarter FE	Y	Y	Y	Y	Y	Y	Y
Excl. Florida	N	Y	N	N	N	N	N
OH and NM recode	N	N	Y	N	N	N	N
Drop OH and NM	N	N	N	Y	N	N	N
Implement after 2005	N	N	N	N	Y	N	N
Implement after 1997	N	N	N	N	N	Y	N
Horwitz legal coding	N	N	N	N	N	N	Y

Notes: This table displays robustness of main results for overdose deaths due to semi-synthetic opioids (T40.2, oxycodone, hydrocodone, morphine, etc.) displayed in Table 4 to alternate specifications. Panel A reports the standard two-way fixed effects model, estimated according to Equation 3. Panel B reports the dynamic ATT for the impact of PDMP introduction, as described in Section 5.2 and Appendix C. County controls include unemployment rate, the percentage of population over 60, and the natural log of population. Robust standard errors, clustered at the state level, in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B3: Opioid overdose deaths (CDC Vital Statistics): Main results and robustness checks

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Two-way fixed effects model						
1(PDMP) _{st}	-0.0713 (0.0453)	-0.0658 (0.0466)	-0.0224 (0.0454)	-0.109 (0.0722)	-0.0876 (0.0553)	-0.877 (0.637)
Panel B: Dynamic ATT						
	-0.108** (0.0439)	-0.104** (0.0442)	-0.0807 (0.0510)		-0.128** (0.0558)	-1.035 (0.637)
Panel C: Dynamic treatment effects						
4 years before	0.0137 (0.0322)	0.0122 (0.0316)	0.0264 (0.0450)	0.0272 (0.0621)	0.0263 (0.0385)	-0.0421 (0.432)
3 years before	-0.0276 (0.0501)	-0.0312 (0.0498)	-0.0216 (0.0583)	-0.0268 (0.0613)	-0.0279 (0.0629)	-0.336 (0.432)
2 years before	0.0133 (0.0228)	0.0107 (0.0228)	0.0155 (0.0224)	0.0249 (0.0321)	0.0210 (0.0265)	0.111 (0.251)
1 year before	0 (.)	0 (.)	0 (.)	0 (.)	0 (.)	0 (.)
Year of implementation	-0.00592 (0.0291)	-0.00433 (0.0296)	0.00543 (0.0246)	-0.0526 (0.0419)	-0.00704 (0.0361)	-0.339 (0.360)
1 year after	-0.0828* (0.0474)	-0.0779 (0.0474)	-0.0819 (0.0554)	-0.149** (0.0730)	-0.0964 (0.0572)	-1.080* (0.632)
2 years after	-0.0995** (0.0436)	-0.0922** (0.0441)	-0.111** (0.0521)	-0.183** (0.0770)	-0.123** (0.0530)	-1.388** (0.633)
3 years after	-0.164*** (0.0571)	-0.154** (0.0571)	-0.173* (0.0893)	-0.228* (0.117)	-0.198*** (0.0703)	-1.645* (0.854)
4+ years after	-0.170** (0.0655)	-0.156** (0.0646)	-0.169* (0.0993)	-0.201 (0.130)	-0.208** (0.0824)	-1.213 (0.986)
Observations	117,856	117,856	117,856	100,864	117,856	117,856
Clusters	38	38	38	38	38	38
Mean	0.62	0.62	0.62	0.62	0.62	5.60
Log-linear model	Y	Y	Y	N	N	N
Poisson model	N	N	N	Y	N	N
arcsinh-linear model	N	N	N	N	Y	N
Linear model	N	N	N	N	N	Y
County controls	Y	N	Y	Y	Y	Y
County FE	Y	Y	Y	Y	Y	Y
Quarter FE	Y	Y	Y	Y	Y	Y
County linear trends	N	N	Y	N	N	N

Notes: This table displays robustness of the main CDC event study specification, displayed in Figure 4 and estimated according to Equation 3, to alternate specifications. Panel A additionally reports the standard two-way fixed effects estimator, estimated using Equation 4. Panel B reports the dynamic ATT for the impact of PDMP introduction, as described in Section 5.2 and Appendix C. County controls include unemployment rate, the percentage of population over 60, and the natural log of population (for log-linear and arcsinh-linear models). The outcome variable for the linear model is the county overdose death rate; the outcome variable for the PPML specification is count of deaths. Robust standard errors, clustered at the state level, in parentheses.

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

Table B4: The Effect of PDMP Introduction on Schedule II Opioids Dispensed (Marketscan): Robustness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
1(PDMP) _{st}	-0.00482*** (0.00160)	-0.00485*** (0.00151)	-0.00522*** (0.00175)	-0.00368*** (0.00126)	-0.00433** (0.00173)	-0.00405** (0.00174)	-0.00455** (0.00169)	-0.00491*** (0.00161)	-0.00272 (0.00228)	-0.00265 (0.00199)
Observations	53,520,024	53,520,024	53,520,024	53,520,024	49,580,421	53,520,024	48,845,666	53,098,616	63,448,213	40,987,179
Clusters	38	38	38	38	37	38	36	36	44	33
Individuals	4,944,277	4,944,277	4,944,277	4,944,277	4,552,407	4,944,277	4,581,977	4,910,548	5,780,901	3,771,915
Individual controls	Y	N	Y	Y	Y	Y	Y	Y	Y	Y
Individual FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Entry-cohort FE	N	N	N	N	N	N	N	N	N	N
Quarter FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
arcsinh-linear model	N	N	Y	N	N	N	N	N	N	N
State linear trends	N	N	N	Y	N	N	N	N	N	N
Excl. Florida	N	N	N	N	Y	N	N	N	N	N
OH and NM recode	N	N	N	N	N	Y	N	N	N	N
Drop OH and NM	N	N	N	N	N	N	Y	N	N	N
Implement after 2005	N	N	N	N	N	N	N	Y	N	N
Implement after 1997	N	N	N	N	N	N	N	N	Y	N
Horwitz legal coding	N	N	N	N	N	N	N	N	N	Y

Notes: This table displays robustness checks corresponding to the baseline Marketscan specification for Schedule II opioid prescribing found in Table 6, Column (1). Individual controls include age, employment type/status [salaried, union, full time, etc.], and industry [oil and gas, manufacturing, etc.]. In Column (6), Ohio and New Mexico are coded with the first NAMSDDL date of implementation in the study period rather than the second. In column (10), I use the legal coding from [Horwitz et al. \(2020\)](#) (date “Modern system operational”), while dropping all states that they indicate had an early PDMP that was legislated and enacted prior to 2003. Robust standard errors, clustered at the state level, in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B5: The Effect of PDMP Introduction on Takeup of Workers' Compensation and Short-Term Disability (Marketscan)

	(1)	(2)	(3)
	Whether worker takes WC (0/1)	Whether WC injury results in days lost (0/1)	Whether worker takes disability (0/1)
1(PDMP) <i>.st</i>	-0.000223 (0.000308)	-0.000131 (0.000708)	0.000170 (0.000450)
Individual Controls	Y	Y	Y
Individual FE	Y	Y	Y
Entry-cohort FE	N	N	N
Quarter FE	Y	Y	Y
Clusters	38	38	38
Observations	53,520,024	4,053,277	53,520,024
Individuals	4,944,277	218,667	4,944,277
Mean	0.005	0.019	0.014

Notes: Dependent variables are binary (0/1) for whether worker makes a workers' compensation (Model (1)) or short-term disability (Model (3)) claim, and binary (0/1) for whether a workers' compensation injury results in days missed (Model (2)), estimated using a linear probability model. Model 1 considers the full Marketscan sample, and is 1 if an enrollee makes a workers' compensation claim under WC in the quarter. Model 2 considers the subsample of enrollees ever injured under workers' compensation, and is 1 if a worker makes a WC claim in the quarter which ever results in days absent. Robust standard errors, clustered at the state level, in parentheses.

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

Appendix C Robustness of Two-Way Fixed Effects Estimators

Recent research has noted that in difference-in-differences specifications with heterogeneity in treatment effects and variation in treatment timing (staggered adoption), the two-way fixed effects (TWFE) estimator can be severely biased away from the parameter of interest (Goodman-Bacon, 2018; Borusyak and Jaravel, 2016; Callaway and Sant’Anna, 2020; Sun and Abraham, 2020; de Chaisemartin and D’Haultfœuille, 2020). As discussed in Goodman-Bacon (2018), the TWFE estimator is comprised of a weighted average of all possible two-by-two difference-in-difference estimators – comparing earlier treated versus later treated as controls, later treated versus earlier treated as controls, treated versus the always-treated, and treated versus the never-treated. The static TWFE estimator (e.g., what I estimate above in Equations 2, 4, and 8) may contain “negative weights” where the estimator is biased in the wrong direction by the use of already-treated units, which may be experiencing dynamic treatment effects, as controls. Sun and Abraham (2020); Callaway and Sant’Anna (2020) and de Chaisemartin and D’Haultfœuille (2020) consider these concerns in the estimation of dynamic treatment effects as well – event studies where leads are used to visualize the placebo time periods for parallel trends, and dynamic treatment effects are visualized with lags, such as estimated above in Equations 1, 3, and 7.⁴⁸

There are reasons to be concerned about this problem in some of my results. My difference-in-differences design consists of a relatively long panel with variation in treatment timing/staggered adoption. As documented in Table B1, later PDMP adopters appear to have stronger implementations than earlier adopters, indicating heterogeneity in treatment effects, and for some of my dynamic visualizations, such as Figure 4, the impact of PDMP introduction appears to be growing over time. In some cases – namely, the opioid overdose deaths results, there is a discrepancy between the magnitude of the dynamic treatment effects estimated using an event-study specification, and the TWFE estimates, which in the opioid overdose deaths results are quite fragile to choice of specification. (The ARCOS results and the Marketscan results are much more stable and the TWFE point estimates typically reflect an approximate average of the dynamic treatment effects.)

To check robustness of my results, I implement the DIDM estimator proposed in de Chaisemartin and D’Haultfœuille (2020) using the Stata package `did_multiplegt`, which estimates all the 2x2 DiD comparisons for groups who switch to becoming treated, against those who remain untreated during this period (i.e., estimates treatment effects using only “clean controls”), for my top-line prescribing and overdose deaths results. I then estimate a specification with 16 quarters of

⁴⁸As noted in Goodman-Bacon (2019), results from Sun and Abraham (2020) suggest that these bias concerns appear more severe for the static rather than dynamic specifications, because the static specifications break down under a wider variety of conditions; a paper that explores this in practice is Cengiz et al. (2019).

leads (placebo estimators) and lags (dynamic treatment effects) to check robustness of the event-study specifications of Equations 1 and 3. Subsequently, I compute a simple average treatment effect from the estimated dynamic effects (lags),⁴⁹ e.g., a “dynamic ATT,” which can be contrasted to the TWFE estimator from Equations 2 and 4.

The dynamic results for ARCOS are presented in Figure C1 (left panel) alongside the original TWFE event study specifications estimated according to Equation 1 (right panel). The shape of the graphs and magnitude of the effect are very similar, indicating the dynamic TWFE specification is robust to the bias problems described above and capture the parameters of interest. The dynamic ATT estimate aggregating these effects for ARCOS is reported in Table B1, Column (11), and is consistent with the TWFE estimates of the baseline specification and robustness checks.^{50,51}

I conduct a similar exercise for the opioid overdose mortality results. The results are presented in Figure C2 (left panel) alongside the TWFE event study specification estimated according to Equation 3 (right panel). Once again, the shape of the graphs and magnitude of the effect are very similar, indicating that the dynamic TWFE specification is robust to the bias problems described above for the overdose deaths models as well.

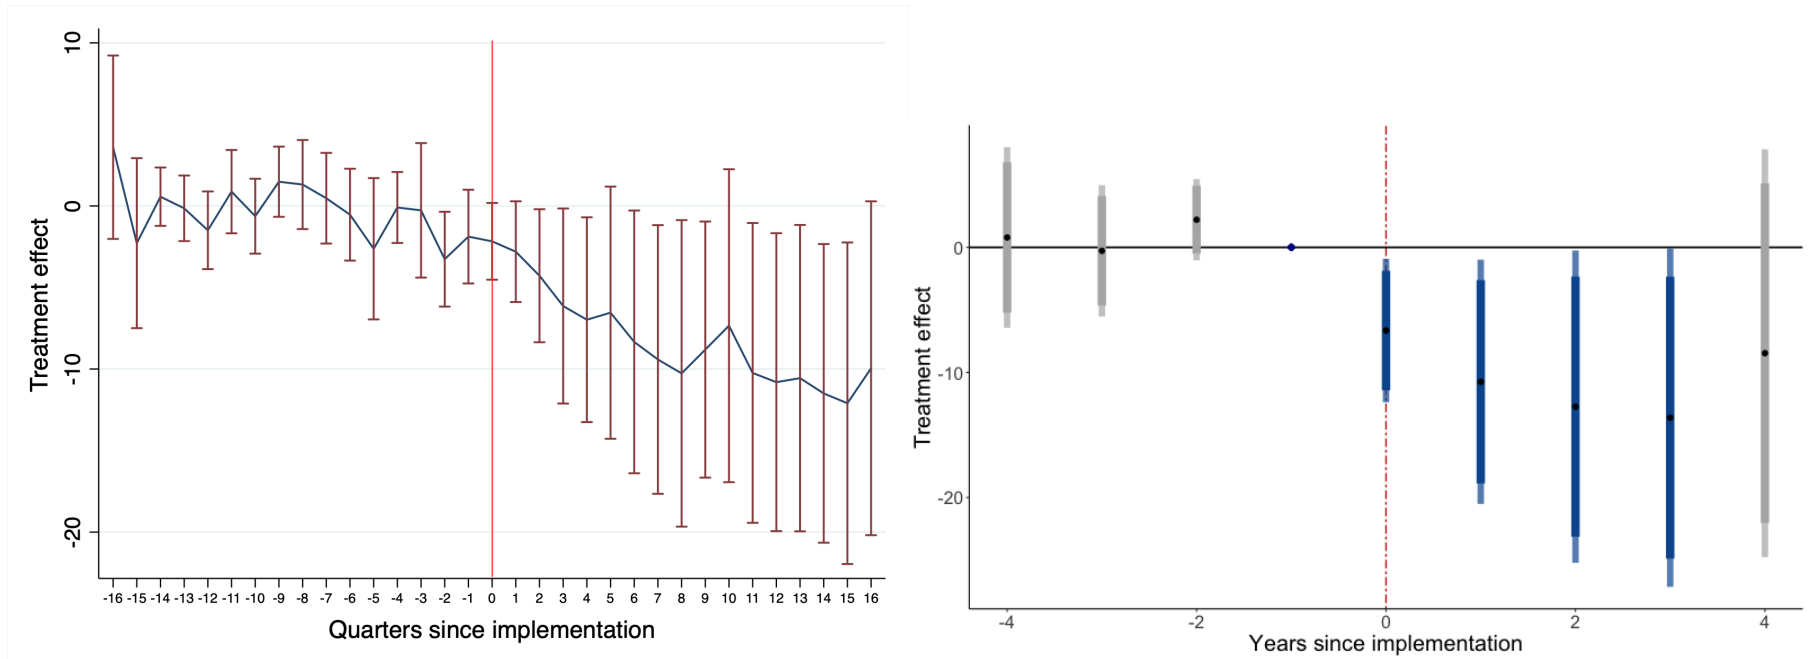
The dynamic ATT is estimated to be -0.108 and is reported in Table B3, Panel (B), Column (1). It is very similar to the simple mean of the dynamic effects in the standard TWFE dynamic specification, as reported in Panel (C), Column (1), which is -0.104. While the TWFE estimator reported in Panel (A), Column (1) is only moderately attenuated relative to the dynamic ATT, the battery of robustness checks presented for the opioid overdose deaths results in Tables B2 and B3, presented across the columns of Panels A and B in both tables, show the TWFE estimator is very fragile to specification choice, while the dynamic ATT estimator is more robust and is closer to the average of the dynamic lags reported in, e.g., Panel C. Throughout the paper, when I report a single point estimate for the overdose death, I report the dynamic ATT rather than the TWFE point estimate; I also rely heavily on the dynamic TWFE event-study specifications, which also appear to be more robust.

⁴⁹Callaway and Sant’Anna (2020) explicitly explore the different ways of aggregating estimates into overall treatment parameters; this approach is averaging across event-time.

⁵⁰This point estimate is very close to a dynamic ATT of -8.6 calculated using the method proposed by Callaway and Sant’Anna (2020), aggregating over event-time.


⁵¹I also conduct a Bacon decomposition of my top-line prescribing results from the ARCOS data, finding the later treated vs. earlier control group’s contribution is nearly identical in magnitude to the overall estimated effect, further alleviating concerns in this specification.

Figure C1: Robustness for dynamic two-way fixed effects specifications robustness (ARCOS): Opioid prescribing



Notes: The left panel depicts 16 placebo estimators and 16 dynamic treatment effects estimated according to the procedure in [de Chaisemartin and D’Haultfœuille \(2020\)](#), implemented using `did_multiplegt`. The right panel reproduces the baseline event study specification depicted in Figure 3. Robust standard errors clustered at state level. In the right panel, thin bars represent 95% confidence intervals and thick bars represent 90% confidence intervals. In the left panel, bars represent 95% confidence intervals.

Figure C2: Robustness for dynamic two-way fixed effects specifications (CDC Vital Statistics): Opioid overdose deaths



Notes: The left panel depicts 16 placebo estimators and 16 dynamic treatment effects estimated according to the procedure in [de Chaisemartin and D’Haultfœuille \(2020\)](#), implemented using `did_multiplegt`. The right panel reproduces the baseline event study specification depicted in Figure 4. Robust standard errors clustered at state level. In the right panel, thin bars represent 95% confidence intervals and thick bars represent 90% confidence intervals. In the left panel, bars represent 95% confidence intervals.

Appendix D Selection into sample: Individual and cohort fixed effects

As discussed above in Section 3, the MarketScan data is not a representative national sample, but rather is a convenience sample obtained from employers (meaning that individuals, whose characteristics are likely to be correlated, enter or leave the sample in groups as their employers join or exit). In Table D1, I consider robustness of the baseline specification to the inclusion of individual fixed effects, which I utilize throughout my analysis. In column (2) I first show that there is a meaningful change in the point estimate of the baseline specification with the full sample if individual FE are not included and instead they are switched with state FE. I then consider the same exercise with a balanced panel in individuals – individuals who are present in my data for 8 years continuously (i.e., 2005-2012) – and estimating the treatment effect with and without individual fixed effects, substituting individual FE with state FE. As shown in Table D1, the point estimates for all these groups are consistently of the same magnitude, which is approximately 2/3 the magnitude of the baseline point estimate for the full sample. That the inclusion of individual FE in this sample does not materially affect the magnitude or significance of the point estimate compared to no individual FE indicates that there are not individual-level unobserved characteristics that might bias the estimate; the individual FE are serving in regressions on the full sample only to deal with changes in the composition of the sample that might bias the estimate. I do not generally utilize this smaller sample of continuous enrollees, but instead opt for the full sample with individual fixed effects, because most individuals are present in my data only for a few years, so there is a considerable cost in terms of power and also representativeness to excluding the modal employee. (It also is likely that individuals employed continuously for many years are different (on observables and unobservables) from employees employed for a shorter duration. They differ on observable demographics – they are on average four years older (46 versus 42), more male (65% versus 60%), are more likely to be in a union (23% versus 17%), and consume fewer opioids on average (60.99 versus 53.03 MME per quarter).)

The ideal specification would be able to utilize include company fixed effects to control for unobservable time invariant unobservables introduced by the shifting sample, but unfortunately company identifiers are not provided in the MarketScan data. Instead I approximate company fixed effects by grouping individuals into 380 cohorts based on the state and year they enter the sample. Robustness to the use of entry-cohort fixed effects rather than individual fixed effects is depicted in Table D1 and Figure D1; results are more precise but very similar. I utilize entry-cohort fixed effects in models where a large number of high-dimensional individual fixed effects are undesirable, namely logistic and Poisson pseudo-maximum likelihood models with many zeros in the dependent variable, and in general do not find that my results meaningfully change when estimated using


entry-cohort FE versus individual FE.

Table D1: The Effect of PDMP Introduction on Schedule II Opioids Distributed: Full Sample and Continuous Enrollees (2005-2012) With and Without Individual FE and Entry-cohort FE (Marketscan)

	(1)	(2)	(3)	(4)	(5)
1(PDMP) <i>_ist</i>	-0.00482*** (0.00160)	-0.00191 (0.00202)	-0.00391 (0.00263)	-0.00323 (0.00288)	-0.00694*** (0.00145)
Sample	Full	Full	Continuously enrolled from 2005-2012	Continuously enrolled from 2005-2012	Full
Individual Controls	Y	Y	Y	Y	Y
Individual FE	Y	N	Y	N	N
Entry-cohort FE	N	N	N	N	Y
State FE	N	Y	N	Y	N
Quarter FE	Y	Y	Y	Y	Y
Observations	53,520,024	53,520,024	6,666,432	6,666,432	53,520,024
Clusters	38	38	38	38	38
Individuals	4,944,277	4,944,277	208,326	208,326	4,944,277
Mean	60.99	60.99	53.03	53.03	60.99

Notes: This table examines the robustness of the point estimates to the exclusion / inclusion of individual fixed effects for the full sample, as well as a restricted subsample of enrollees that provide a balanced panel in 8 years, as well as to entry-cohort fixed effects as alternate controls to individual fixed effects. Dependent variable is $\ln(\text{MME Sch. 2} + 1)$. Means are in levels (MME). Robust standard errors, clustered at the state level, in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure D1: The Effect of PDMP Introduction on Schedule II Opioids Distributed using Individual versus Entry-cohort fixed effects (MarketScan)



Notes: The left panel depicts the baseline specification with 4,944,277 individual FE; the right panel with 380 entry-cohort FE. Robust standard errors clustered at state level. Thin bars represent 95% confidence intervals and thick bars represent 90% confidence intervals.

Appendix E Selecting the chronic pain cohort

A large number of medical conditions can cause chronic pain, ranging from lower back injury to cancer pain to migraines and nerve pain, and no one diagnosis or procedure code is used consistently in claims data to indicate an individual that is suffering from chronic pain. In order to identify a sample of likely chronic pain sufferers, I employ a predictive machine learning algorithm that is able to flexibly identify associations between a large number of clinical covariates, which can be extracted from the Marketscan medical claims and work data used in this study, and the likelihood of observing a chronic opioid prescription. While not all pain sufferers are prescribed a long-term opioid prescription, as has been demonstrated throughout this paper, I assume that the covariates most frequently observed as correlating with such a long-run opioid prescription are the covariates that are most likely to select chronic pain patients. Below and in [Appendix F I](#) characterize the common groupings of care that the chronic pain cohort receives. This algorithmic approach can be contrasted with simpler methods of selecting pain patients based on the presence of a list of diagnosis codes, such as is described in [Tian, Zlateva and Anderson \(2013\)](#). For my purposes, generating a chronic pain cohort with high rates of opioid usage and indicators of serious chronic pain conditions, the algorithmic approach is better-performing. Specifically, selecting based on broad pain codes (such as the approximately 100 in [Tian, Zlateva and Anderson](#)) generates a large sample of over a million individuals with only moderately elevated opioid prescribing and fewer indicators of serious chronic pain.

Appendix E.1 Data

The data used to train the machine learning model are extracted from commercial health insurance claims data: Marketscan covering years 2005-2012.

Appendix E.2 Labels

The outcome variable of interest is the existence of a chronic opioid prescription, encoded as a binary classifier. I use a custom episode grouper to link related claims across time, and consider a person to have a chronic opioid prescription if they have a prescription lasting more than 90 days.

Appendix E.3 Features

Covariates (model features) are extracted from claims data records, and indicators of clinical experience are transformed to a panel data structure with observations at the individual-month level. Covariates used include sex, 5-year age bins (Age 26-30, etc.), dummy variables for industry

of employment, and a large panel of clinical variables, including dummy indicators for 3-digit ICD-9 codes, facility types, and provider types. In total, 1,371 features were included in the prediction model.

Appendix E.4 Algorithm

I utilize a binary classification machine learning algorithm, gradient boosted decision trees, implemented in Python by the XGBoost library. The algorithm, which is an ensemble method, fits a large number of trees sequentially (each tree allowing for a non-linear relationship structure between covariates) with iterative improvements each time (Friedman, Hastie and Tibshirani, 2009). The resulting model generates a score at the individual-month level that reflects the machine-determined probability of a chronic opioid therapy prescription in that month; in a given month, if an individual has features that are correlated in the broader sample with a chronic opioid prescription, they will receive a higher score than individuals who do not. I determine an individual's overall likelihood of being a chronic pain patient according to their *maximum* machine-predicted score over their period in the sample, and I thus select observations into or out of the sample at the *individual level*.

Appendix E.5 Cohort Selection

A final sample is constructed from patients who have eligibility for either workers compensation, short term disability, or both, throughout their time in the sample; this is so as to not bias the work results. I then take predicted pain scores from the machine learning algorithm, dropping any enrollees used for training, and augment the scores of employees ever injured enough to miss days of work under workers' compensation, replacing their score with a 1 and ensuring inclusion in the sample. From this score dataset I select the top decile.

The augmentation procedure for workers compensation injured individuals accounts for around 38,000 individuals in the chronic pain patient sample, out of 340,660. I include these individuals because the nature of workers' compensation injuries means that they almost always are accompanied by pain: about 70% of workers injured enough to miss days of work received an opioid prescription.⁵² Further, because medical claims directly related to the workers compensation injury are not observable in detail in the Marketscan databases (the medical claims are paid directly by workers compensation and costs are only reported in aggregate) I will not reliably be able to observe the target classifier of having a chronic opioid prescription.

⁵²Examples of injuries severe enough to miss work in my dataset include "slip and fall", "caught in/between machinery", "continuous trauma", "strain or injury by lifting", "struck/injured by falling or flying object", etc.

Appendix E.6 Model results

Table E1 displays the results of a test for whether the pain scores are endogenous to PDMP introduction; they are not. Table B5, above, also shows that WC claiming behavior is not endogenous to PDMP introduction, so the addition of individuals injured badly enough on workers compensation to miss work is also not endogenous to PDMP introduction. Taken together, these results indicate that even though prescribing behavior and the treatment of chronic pain has changed throughout the sample period, the method used to construct the pain cohort is time-invariant enough for selecting a consistent cohort of individuals, and individuals in the chronic pain sample are not being selected into the sample differently in later years versus earlier years.

Table E1: Endogeneity in Predicted Enrollee Sample (Marketscan)

	(1) Enrollee score in predictive model
$1(\text{PDMP})_{ist}$	-0.000217 (0.000916)
Controls:	
Individual Controls	Y
Individual FE	N
Entry-cohort FE	Y
Quarter FE	Y
State FE	N
Clusters	38
Observations	37,574,510

Notes: This table displays the relationship between the introduction of a PDMP and an enrollee's score predicting their likelihood of being a chronic pain patient. Robust standard errors, clustered at the state level, in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Characteristics of the chronic pain cohort are reported in Table E2. Chronic pain patients are older, more female, and more likely to be blue-collar (working in union and hourly employment rather than non-union or salaried; working in manufacturing rather than finance, insurance, or real estate).

Table E2: Characteristics of chronic pain sample (Marketscan)

	(1) Full sample excluding chronic pain	(2) Chronic pain sample
Female	0.399	0.442
Age	41.8	46.0
<i>Age bins:</i>		
17-28	0.148	0.056
29-40	0.308	0.233
41-52	0.349	0.419
53-64	0.195	0.292
<i>Geography:</i>		
Northeast	0.172	0.171
North Central	0.221	0.252
South	0.351	0.335
West	0.255	0.243
Lives in Metro Area	0.907	0.908
<i>Employment status:</i>		
Union	0.155	0.242
Non-union	0.613	0.557
Salaried	0.432	0.356
Hourly	0.334	0.433
<i>Industry:</i>		
Oil & Gas Extraction, Mining	0.017	0.006
Manufacturing, Durable Goods	0.265	0.373
Manufacturing, Nondurable Goods	0.136	0.132
Transportation, Communications, Utilities	0.173	0.196
Retail Trade	0.026	0.012
Finance, Insurance, Real Estate	0.232	0.159
Services	0.127	0.076
Individuals	4,603,617	340,660

Notes: Reports means of demographic indicators for individuals in the chronic pain sample (Column 2), compared to individuals in the full sample excluding the chronic pain sample (Column 1).

Appendix F Grouping and characterizing episodes of care

I characterize the experience of the cohort of chronic pain patients using latent Dirichlet allocation (LDA) (Blei, Ng and Jordan, 2003). LDA is an unsupervised natural language processing method that probabilistically models unobserved underlying topics as generating observed collections of words.

I first segment episodes of care for the chronic pain sample by grouping all the claims that occur into episodes, which I define as any period of time in which there is at least one claim in each continuous 3-month period. I extract all codes from each claim of each episode, including all 3-digit ICD-9 diagnosis codes, major diagnostic category codes, inpatient discharge statuses, CPT procedure codes, procedure groups, provider types, service categories, places of service, revenue codes, etc. I combine all episode-associated codes into one “bag of words”, and each “document” supplied to the algorithm is a collection of all the words (claims) associated with an episode of care. I implement an LDA model for the identification of 15 topics using the LDA algorithm implementation in Python’s `sklearn`, and visualize and interpret the extracted topics using the `pyLDAvis` package.⁵³

Topic models such as LDA are models of how *latent topics* generate observed *documents* containing *words*. Each document is modeled as a probabilistic mixture of topics; each of those latent topics are associated with distributions of words, and the topic model identifies the latent topics and their associated word distributions. By applying LDA to patient episodes, I model the *patient episode record* (“documents”) comprising claims codes (“words”) as arising from latent *health conditions* (“topics”). Thus each patient-episode may have the observed claims generated by one or several underlying health conditions.

As can be seen in Figures F1–F5, the LDA algorithm easily identifies interpretable episodes of patient care. It identifies episodes of care relating directly to the treatment of pain conditions (e.g., surgery, and outpatient physical therapy, for musculoskeletal and joint pain); episodes of care that may have pain as a side effect (e.g., cancer); and unrelated episodes of care that occur at other points during the patients’ time in the sample. My one-line characterization of the contents of each of the 15 medical topics are reported in Table F1. I take the further step of identifying topics that identify pain episodes specifically, as well as other types of care of interest, like routine medical care, ER visits, and other serious medical problems which are not pain conditions but might require opioid pain management in addition to other treatment (e.g., cancer). I report those in Column (2). As discussed above in the text, I consider outcomes of interest - overall inpatient and outpatient spending, and missed work, in each of these kinds of patient episodes, finding in Table F2 that

⁵³The `pyLDAvis` interactive visualizations of all fifteen groups are viewable on https://angelakilby.com/html/LDA-pain_patients.html.

PDMP introduction and opioid prescribing reduction has an observed impact specifically on work function and overall costs of pain management care for patients experiencing pain episodes, e.g., episodes that are characterized by the graphics in Figures F1 and F2.

Table F1: Descriptions of 15 topics identified by LDA topic model

Topic	Description	Table F2 grouping
1	Vision	
2	Inpatient musculoskeletal surgery	Chronic pain treatment episodes
3	Female-related mammograms and surgeries	
4	Routine office visits	Routine medical care
5	Routine female-related	Routine medical care
6	Cancer	Other serious medical problems
7	Office visits for respiratory infections	
8	Outpatient musculoskeletal, back, and spinal pain treatment	Chronic pain treatment episodes
9	Cardiovascular/heart disease	Other serious medical problems
10	Post-hospitalization aftercare	
11	Labwork, preventative, diagnostics	
12	ER visits	ER visits
13	Inpatient stay	Other serious medical problems
14	Mental health and depression	
15	Neurology, sleep	

Figure F1: Topic modeling example: Inpatient musculoskeletal surgery



Figure F2: Topic modeling example: Outpatient musculoskeletal, back, and spinal pain treatment




Figure F3: Topic modeling example: ER visits




Figure F4: Topic modeling example: Routine office visits

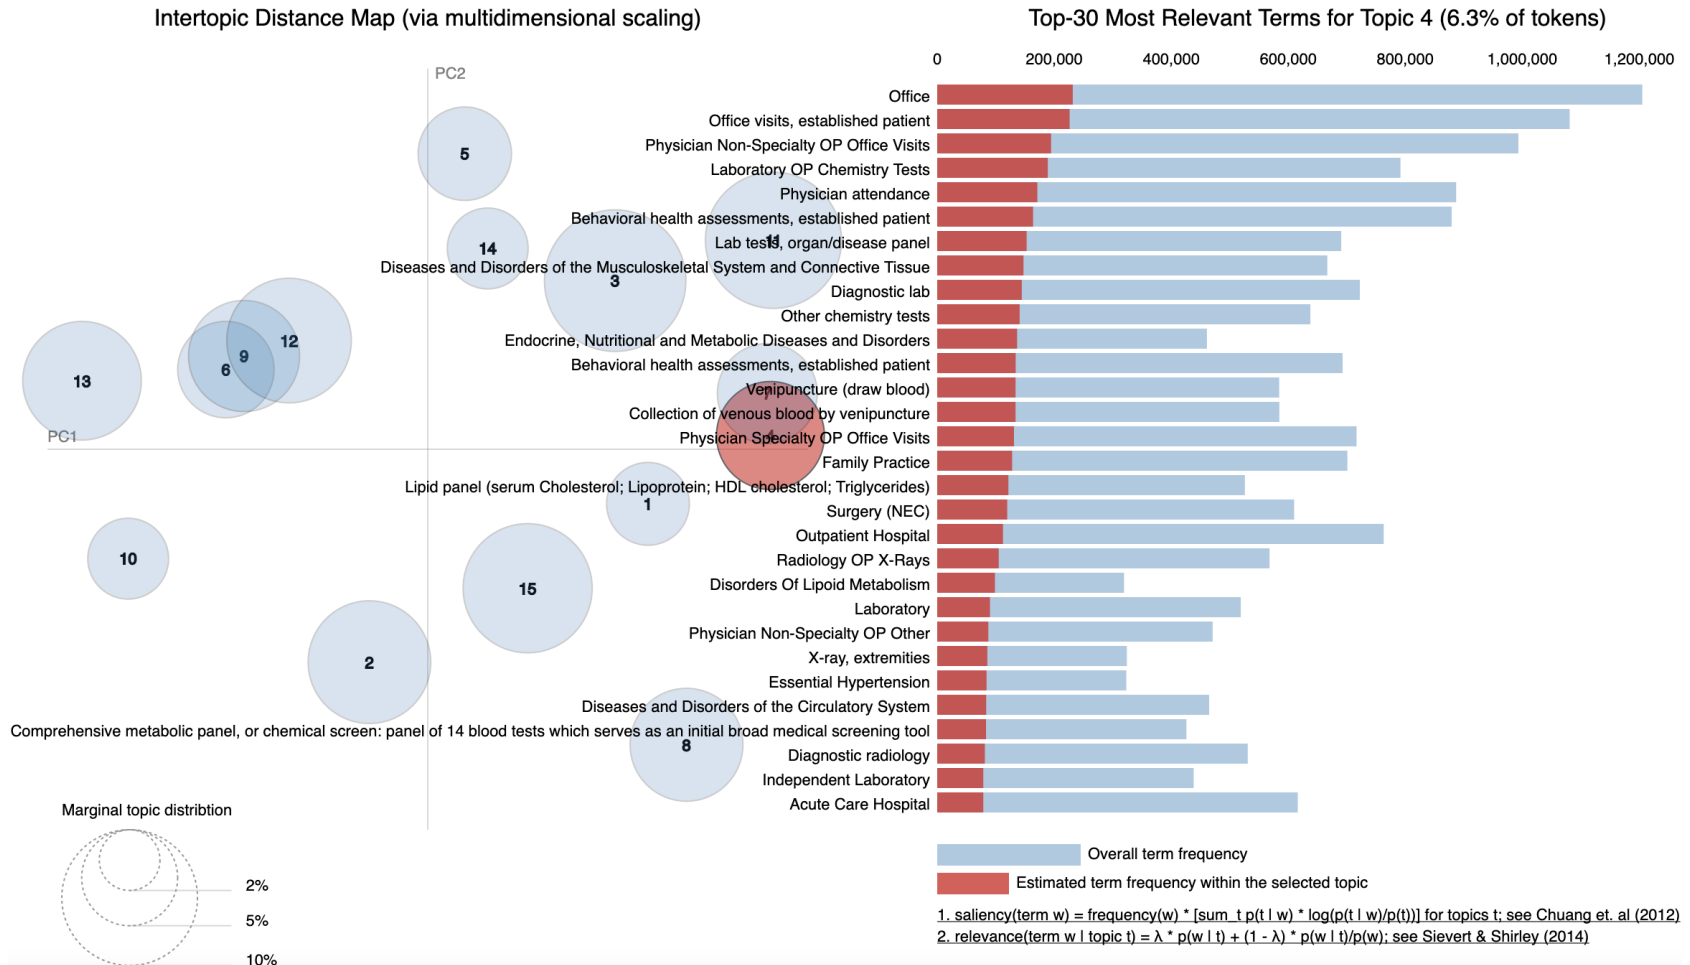


Figure F5: Topic modeling example: Cancer

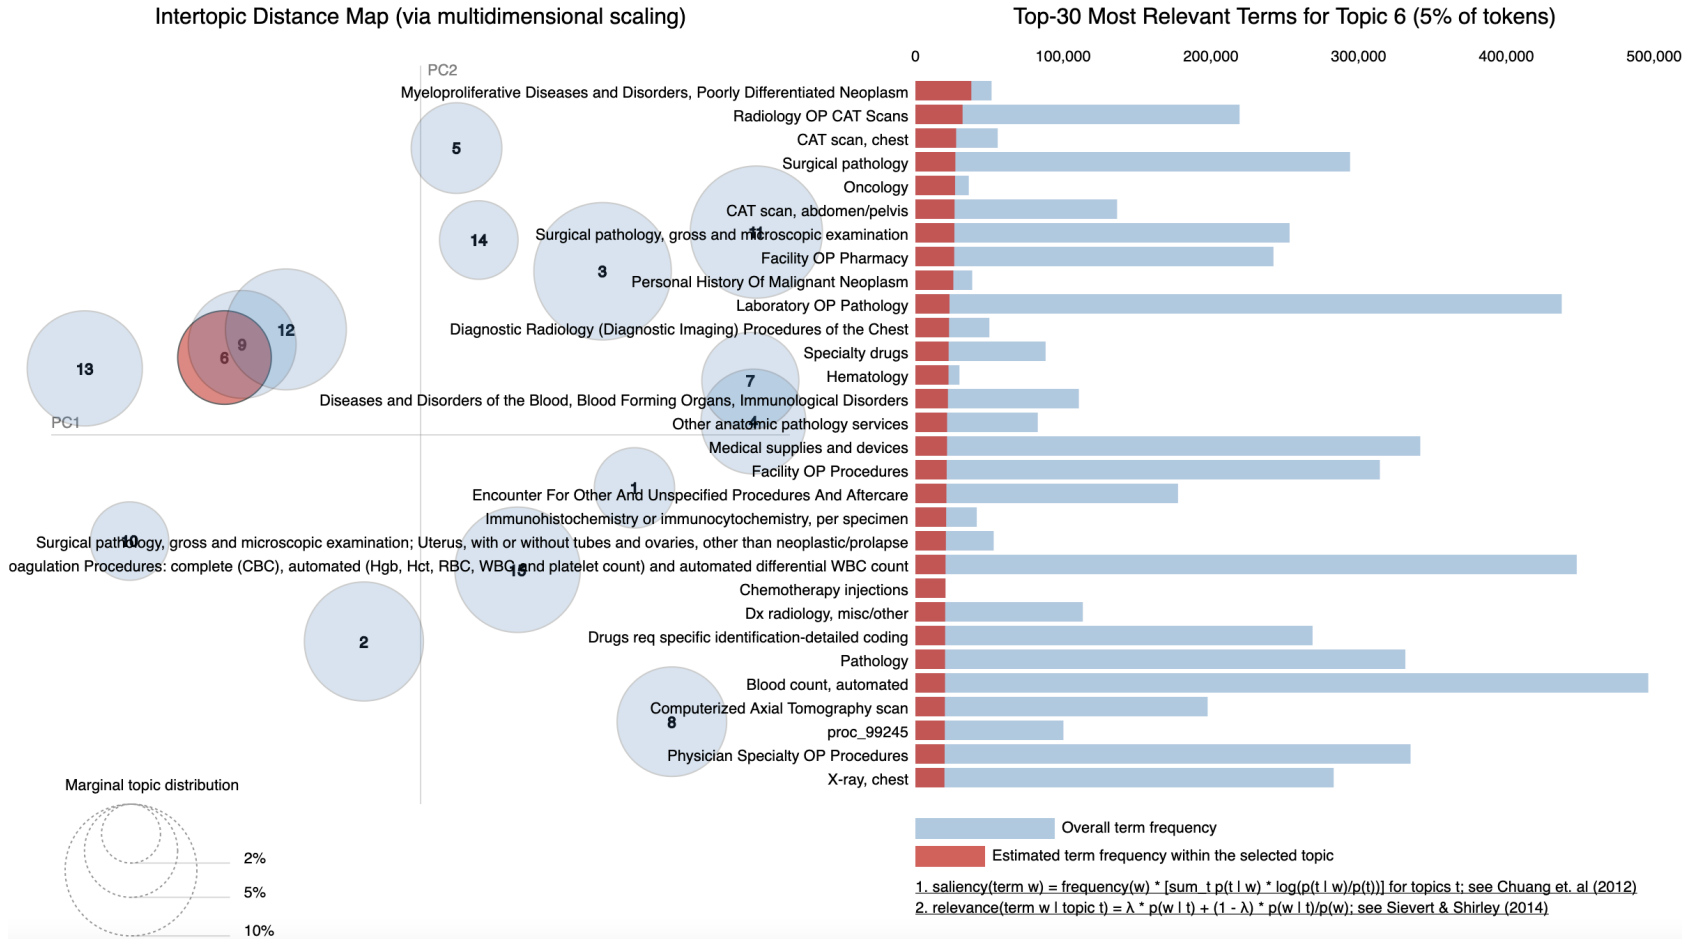


Table F2: Impact of PDMP introduction on outcomes during episodes of care identified by the LDA algorithm (Marketscan)

	(1)	(2)	(3)	(4)
$1(\text{PDMP})_{st}$	Chronic pain treatment episodes	ER visits	Routine medical care	Other serious medical problems
$\log(\text{spending}+1)$	0.0223** (0.00959)	0.0139 (0.0134)	0.00967 (0.0206)	0.0196 (0.0126)
$\log(\text{days absent} + 1)$	0.0229*** (0.00764)	0.00434 (0.0139)	0.00986 (0.0111)	0.0202** (0.00837)
Observations	956,389	478,960	505,714	1,051,176

Notes: Table reports the impact of PDMP introduction on inpatient and outpatient spending and days absent during periods that the topic modeling algorithm has identified as being episodes where the patient is being treated for chronic pain. Table F1 reports which topic models are included in which of the above four categories.