The Welfare Cost of Prescription Drug Monitoring Programs: A Sufficient Statistics Approach

APEP Autonomous Research* @SocialCatalystLab

March 2026

Abstract

Must-access Prescription Drug Monitoring Programs require physicians to query a state database before prescribing opioids. We derive sufficient statistics for the welfare effect of prescribing regulation under three behavioral models of addiction—rational, present-biased, and cue-triggered—showing that the welfare sign depends on a single parameter: consumer bias β^*. Using staggered adoption across 36 states (2012–2019) and a Medicare Part D panel spanning 2013–2023, we estimate the causal effect on opioid prescribing via the Callaway and Sant'Anna (2021) estimator. The aggregate treatment effect is -0.070 percentage points (SE: 0.102), a 1.2% reduction from baseline. At baseline targeting ($\lambda = 0.70$), welfare is negative unless present bias is very strong ($\beta < 0.37$). Better-targeted PDMPs ($\lambda = 0.30$) raise the threshold to $\beta^* = 0.90$. The welfare question depends jointly on the behavioral model and on how well the policy targets addiction-risk prescriptions.

JEL Codes: I12, I18, D60, H51

Keywords: opioids, PDMP, welfare analysis, sufficient statistics, addiction, prescribing regulation

*Autonomous Policy Evaluation Project. Code and data: https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0488. Total execution time: N/A (cumulative: N/A). All errors are our own. Correspondence: scl@econ.uzh.ch

1. Introduction

Every opioid prescription prevented by regulation is either a future addiction averted or a patient left in pain. Thirty-six states now require physicians to check a database before writing an opioid prescription—and these must-access mandates have reduced prescribing by 8 to 12 percent (Buchmueller and Carey, 2018). The reduction is celebrated as evidence that Prescription Drug Monitoring Programs "work." But "work" is a welfare claim, not a statistical one.

The welfare sign depends on which effect dominates—addiction averted or pain inflicted—and that depends, in turn, on two empirical quantities: how biased are the patients whose prescriptions were prevented, and how many of them were addiction-risk patients versus legitimate pain patients? A blunt mandate that mostly prevents prescriptions to chronic pain patients (high targeting ratio λ) destroys welfare even if consumers are substantially biased. A well-targeted mandate that selectively reduces addiction-risk prescriptions improves welfare under far weaker behavioral assumptions.

This paper provides the first welfare analysis of prescription drug monitoring programs. We derive the sufficient statistics for the welfare effect of a supply-side prescribing regulation under addiction, estimate the key statistic—the causal effect of must-access mandates on prescribing—from staggered state adoption, and present welfare bounds under three canonical behavioral models of addiction. The main result is a threshold theorem: the welfare sign depends on a critical value β^* that is itself a function of how well the PDMP targets addiction-risk prescriptions. Under current hassle-cost-driven designs ($\lambda \approx 0.70$), $\beta^* \approx 0.37$—below most experimental estimates of present bias, implying that PDMPs are welfare-reducing. Better-targeted designs would raise β^* substantially.

Our conceptual framework extends the sufficient statistics approach of Chetty (2009) and Allcott et al. (2019) to gatekeeper-mediated regulation of an addictive good. The standard sufficient statistics welfare formula for a sin tax decomposes welfare into internality correction, externality correction, and fiscal effects. Our setting differs in three ways. First, the regulation operates through a supply-side gatekeeper (the prescriber) rather than through prices, introducing a physician agency wedge into the welfare formula. Second, the good is addictive, so the internality is dynamic: today's prescription generates future addiction costs that a present-biased consumer underweights. Third, the regulated quantity is a medical treatment with direct therapeutic value, so the welfare cost of restricting access is a health cost—forgone pain management—rather than a consumption loss.

We formalize these features in a two-type model where patients differ in addiction risk ($\theta \in \{L, A\}$) and may exhibit quasi-hyperbolic discounting with parameter $\beta \in [0, 1]$ (Laibson,

1997). A social planner chooses regulatory stringency τ—the PDMP mandate—which reduces prescribing to both types. The welfare formula decomposes into five sufficient statistics: the per-patient internality γ, the externality e, the physician agency wedge ϕ, the pain management value v_L, and the targeting parameter λ (the share of prescribing reduction falling on legitimate pain patients). These statistics can be calibrated from the existing literature on addiction costs, overdose externalities, and pain management.

Proposition 1 establishes that welfare from the PDMP is:

$$\frac{dW}{d\tau} = \left[e + (\gamma + \bar{\phi})(1-\lambda) - v_L\lambda\right] \cdot \left(-\frac{d\bar{Q}}{d\tau}\right) - C'(\tau)$$

where λ is the share of the prescribing reduction falling on legitimate pain patients and $(1-\lambda)$ is the share falling on addiction-risk patients. **Corollary 1** provides the sign condition: welfare is positive if and only if $\beta < \beta^*$, where the critical threshold β^* depends on both the behavioral parameters and the policy's targeting efficiency λ.

For the empirical component, we exploit the staggered adoption of must-access PDMP mandates across 36 states between 2012 and 2019. Our data combine CMS Medicare Part D Opioid Prescribing by Geography files (2013–2023) with the RAND OPTIC policy database, CDC provisional overdose mortality data, and Census/FRED economic controls. We estimate the causal effect of must-access mandates using the Callaway and Sant'Anna (2021) doubly robust estimator, which is robust to treatment effect heterogeneity across adoption cohorts. Our main sample restricts to the 30 states adopting mandates in 2014 or later plus 15 never-treated states, providing 495 state-year observations with at least one year of pre-treatment data for every cohort.

The CS-DiD aggregate treatment effect on the opioid prescribing rate is -0.070 percentage points (SE: 0.102), a 1.2 percent reduction from the baseline mean of 6.03 percent. The estimate is economically modest and statistically imprecise—the 95 percent confidence interval includes both a 4.5 percent reduction and a 2.2 percent increase. Traditional two-way fixed effects yield a similar estimate of -0.063 percentage points (SE: 0.086). The CS-DiD effect on the share of prescribers writing any opioid prescription is $+0.009$ (SE: 0.028)—positive but statistically indistinguishable from zero. The only marginally significant result is a small increase in the long-acting opioid share ($+0.0024$, SE: 0.0014), suggesting composition effects in prescribing patterns. Pre-trends are excellent: a joint test of four pre-treatment event-study coefficients yields $\chi^2(4) = 0.54$ ($p = 0.97$), providing strong support for the parallel trends assumption.

These results are consistent with the growing literature finding that must-access PDMP mandates have modest effects on opioid prescribing (Kaestner and Ziedan, 2019; Horwitz et

al., 2021; Dave et al., 2021), particularly in the Medicare population where prescribing was already declining from other forces. The prescribing reduction is considerably smaller than the 8–12 percent found by Buchmueller and Carey (2018) using earlier data (2006–2014), consistent with attenuation as later-adopting states had lower baseline prescribing rates and as secular trends in opioid awareness reduced prescribing nationwide.

The welfare calibration is the paper's central contribution. Feeding our estimated prescribing reduction through the sufficient statistics formula, we find:

- Under **rational addiction** ($\beta = 1$): the externality correction ($500) falls far short of the pain cost ($5,250 at $\lambda = 0.70$). Net welfare is −$4,750 per prevented prescription.
- Under **moderate present bias** ($\beta = 0.7$): the internality correction applies only to the 30% of prevented prescriptions falling on addiction-risk patients. Net welfare is still negative at −$2,500.
- Under **cue-triggered addiction** ($\beta = 0$): the full internality correction dominates. Net welfare is +$2,750.

The critical threshold at baseline targeting is $\beta^* \approx 0.37$—well below most experimental estimates of present bias in substance users ($\beta \approx 0.4$–0.7). But β^* is highly sensitive to targeting: at $\lambda = 0.30$ (a well-targeted PDMP), β^* rises to 0.90. The paper's central finding is that the welfare case for PDMPs depends as much on how well the policy targets addiction-risk prescriptions as on the behavioral model of addiction.

This paper contributes to four literatures. First, we contribute to the literature on PDMPs and opioid prescribing (Buchmueller and Carey, 2018; Kaestner and Ziedan, 2019; Horwitz et al., 2021; Dave et al., 2021; Meinhofer, 2018; Maclean et al., 2022). Existing work estimates prescribing effects but does not evaluate welfare. We provide the first welfare analysis by embedding the prescribing estimate in a structural model of addiction.

Second, we contribute to the sufficient statistics approach to welfare analysis (Chetty, 2009; Allcott et al., 2019; Hendren and Sprung-Keyser, 2020). The prior literature applies sufficient statistics to sin taxes (cigarettes, sugar), disability insurance, and health insurance. We extend the framework to supply-side prescribing regulation, introducing the physician agency wedge as a new channel and the addictive good as a setting where the internality is dynamic and bimodal.

Third, we contribute to the behavioral economics of addiction (Becker and Murphy, 1988; Gruber, 2001; Bernheim and Rangel, 2004; O'Donoghue and Rabin, 2006; Gruber and Köszegi, 2001). The theoretical literature proposes competing models with sharply different welfare implications, but empirical work has not mapped these models to specific policy evaluations.

We show how the sufficient statistics formula can be parameterized under each model, making the welfare implications of competing addiction theories transparent.

Fourth, we contribute to the literature on the opioid crisis (Case and Deaton, 2015; Currie et al., 2019; Ruhm, 2019; Volkow and McLellan, 2016; Kolodny et al., 2015; Alpert et al., 2022). The crisis literature has documented causes, consequences, and policy effects, but welfare evaluation remains rare. Mulligan (2020) models nonconvex budget sets in opioid markets but does not quantify welfare. Schnell (2017) documents physician heterogeneity in prescribing but does not connect it to welfare. We provide a tractable framework that maps the extensive empirical knowledge about opioid prescribing into welfare statements.

The remainder of the paper proceeds as follows. Section 2 develops the conceptual framework and derives the welfare formula. Section 3 describes the institutional background of PDMP mandates. Section 4 describes the data. Section 5 presents the empirical strategy. Section 6 reports the main results. Section 7 calibrates the welfare formula. Section 8 presents robustness checks. Section 9 discusses limitations and extensions. Section 10 concludes.

2. Conceptual Framework

We develop a model of prescribing regulation for an addictive good. The model nests three canonical theories of addiction—rational, present-biased, and cue-triggered—within a single sufficient statistics framework, following the approach of Chetty (2009) and Allcott et al. (2019).

2.1 Setup

Consider a continuum of patients indexed by type $\theta \in \{L, A\}$, where L denotes legitimate pain patients and A denotes addiction-risk patients. The population share of type A is π. Each patient has a physician who prescribes on the patient's behalf.

Time is discrete with periods $t \in \{0, 1, 2, \ldots\}$. In period 0, the physician decides whether to prescribe an opioid. If prescribed, the patient receives immediate utility $u(\theta)$ from pain relief. In subsequent periods, addiction-risk patients face an expected addiction cost stream with undiscounted present value $K > 0$ (evaluated from period 1 onward), encompassing physical dependence, tolerance, escalation, and overdose risk. At the prescribing decision in period 0, the normatively correct present value of these costs is δK, where δ is the exponential discount factor. Legitimate pain patients face no addiction cost.

Assumption 1 (Quasi-Hyperbolic Discounting)**.** *Patients discount future costs using the* (β, δ) *framework of Laibson (1997), where* $\beta \in [0, 1]$ *captures present bias and* δ *is the*

standard exponential discount factor. The perceived cost of future addiction at the time of the prescribing decision is $\beta\delta K$ rather than the normatively correct δK.

Under Assumption 1, the parameter β nests the three addiction models:

- **Rational addiction** (Becker and Murphy, 1988): $\beta = 1$. The patient fully internalizes future addiction costs.
- **Present-biased addiction** (Gruber, 2001; Gruber and Köszegi, 2001): $\beta \in (0, 1)$. The patient underweights future costs by factor β.
- **Cue-triggered addiction** (Bernheim and Rangel, 2004): $\beta = 0$. In the cue-triggered state, the patient places zero weight on future costs.

Definition 1 (Internality)**.** *The internality for a type-A patient with bias parameter β is:*

$$\gamma(\beta) = (1 - \beta) \cdot \delta K \tag{1}$$

This is the wedge between the patient's perceived cost and the true cost of addiction at the time of the prescribing decision.

The share of patients who are addiction-risk types is π. In the welfare formula, π does not appear explicitly because Q_A and Q_L are aggregate quantities; π is implicitly captured by the targeting parameter λ (defined below).

2.2 Physician Agency

Physicians prescribe on behalf of patients but may deviate from the patient's true optimum. We model the physician agency wedge $\phi_j \geq 0$ for physician j as an additive utility wedge in the physician's prescribing rule: the physician prescribes as if the patient's marginal benefit is $v_A + \phi_j$ rather than v_A, leading to the marginal condition $v_A + \phi_j = \beta\delta K$ (the perceived cost). Following Schnell (2017), this wedge arises from time pressure, patient demand accommodation, incomplete information about addiction risk, and pharmaceutical marketing.

Assumption 2 (Physician Agency)**.** *Let $\bar{\phi} = \mathbb{E}[\phi_j | j$ prescribes opioids$]$ denote the average agency wedge among opioid prescribers. We assume $\bar{\phi} \geq 0$: physicians weakly overprescribe relative to the patient's long-run optimum.*

The agency wedge is a sufficient statistic for all supply-side distortions. In utility terms, ϕ_j measures the dollar-equivalent "taste for prescribing" that leads the physician to prescribe

beyond the patient's long-run optimum. Three sources generate positive $\bar{\phi}$ in the opioid context. First, *time pressure*: an office visit lasts 15 minutes on average, and writing a prescription is faster than explaining non-opioid pain management alternatives. Second, *demand accommodation*: patients request opioids, and physicians face social costs from refusing. Third, *incomplete information*: the physician observes the patient's current pain but not the patient's addiction risk type θ, leading to overprescribing on the margin where the physician is uncertain.

2.3 The Planner's Problem

A social planner chooses regulatory stringency $\tau \geq 0$, where $\tau = 0$ is no regulation and higher τ represents more stringent prescribing requirements (the PDMP mandate). The regulation reduces prescribing for both patient types: total prescribing $\bar{Q}(\tau)$ is decreasing in τ, with $Q_L(\tau)$ and $Q_A(\tau)$ denoting prescribing to legitimate pain and addiction-risk patients, respectively.

Each prescription generates an externality $e > 0$ capturing overdose costs borne by others (emergency response, family disruption, workplace productivity, healthcare costs to insurers). The planner maximizes welfare $W(\tau)$, evaluated using the patient's long-run preferences (i.e., using δK rather than $\beta\delta K$):

$$W(\tau) = \underbrace{v_L \cdot Q_L(\tau) + v_A \cdot Q_A(\tau)}_{\text{pain relief}} - \underbrace{\delta K \cdot Q_A(\tau)}_{\text{addiction cost (true)}} - \underbrace{e \cdot \bar{Q}(\tau)}_{\text{externality}} - \underbrace{C(\tau)}_{\text{admin cost}} \tag{2}$$

Here $Q_L(\tau)$ and $Q_A(\tau)$ are *aggregate* prescribing quantities (total prescriptions across all patients of each type), so $\bar{Q}(\tau) = Q_L(\tau) + Q_A(\tau)$. The pain relief term captures the direct health benefit: v_L per prescription for legitimate pain patients, v_A for addiction-risk patients. The addiction cost is evaluated at the true (not perceived) present value δK. The externality e is proportional to total prescribing. The administrative cost $C(\tau)$ captures the direct cost of operating the PDMP system (database maintenance, prescriber time, licensing).

Three features distinguish this planner's problem from the standard sin tax problem in Allcott et al. (2019). First, the instrument is a quantity restriction (mandated information acquisition) rather than a price (tax). The PDMP does not raise the price of opioids; it raises the time cost of prescribing them. Second, the regulation is mediated by a gatekeeper (the physician) whose behavior may be misaligned with the patient's long-run interest. Third, the regulated good has direct therapeutic value—it relieves pain—creating a first-order welfare cost of restricting access that does not arise for cigarettes or sugary drinks.

2.4 Deriving the Welfare Formula

Differentiating equation (2) with respect to τ yields the marginal welfare effect of increasing regulatory stringency.

Proposition 1 (Sufficient Statistics Welfare Formula)**.** *Under Assumptions 1–2, define* $\lambda \equiv (-dQ_L/d\tau)/(-d\bar{Q}/d\tau) \in [0,1]$ *as the share of the prescribing reduction borne by legitimate pain patients. The marginal welfare effect of the PDMP mandate is:*

$$\frac{dW}{d\tau} = \underbrace{\left[e + (\gamma + \bar{\phi})(1-\lambda) - v_L\lambda\right]}_{\text{net benefit per prevented Rx}} \cdot \underbrace{\left(-\frac{d\bar{Q}}{d\tau}\right)}_{\substack{\text{prescribing} \\ \text{reduction}}} - \underbrace{C'(\tau)}_{\text{admin cost}} \tag{3}$$

where:

- $\gamma = (1-\beta)\delta K$ *is the per-patient internality for addiction-risk types*
- e *is the externality per prescription (overdose deaths, healthcare costs, diversion)*
- $\bar{\phi}$ *is the average physician agency wedge per prescription*
- v_L *is the marginal value of pain management for legitimate pain patients*
- $(1-\lambda)$ *is the share of the reduction falling on addiction-risk patients (the "targeting" parameter)*
- $C'(\tau)$ *is the marginal administrative cost*

Proof. Differentiate $W(\tau)$ in equation (2). Since Q_L and Q_A are aggregate quantities with $\bar{Q} = Q_L + Q_A$:

$$\frac{dW}{d\tau} = \underbrace{v_L \cdot \frac{dQ_L}{d\tau}}_{\text{pain relief lost}} + \underbrace{(v_A - \delta K) \cdot \frac{dQ_A}{d\tau}}_{\text{net effect on A-types}} - \underbrace{e \cdot \frac{d\bar{Q}}{d\tau}}_{\text{externality saved}} - C'(\tau) \tag{4}$$

For an at-risk patient at the unregulated margin, the physician's prescribing condition is $v_A + \phi = \beta\delta K$ (the physician prescribes until the combined benefit equals perceived cost), so $v_A - \delta K = -(1-\beta)\delta K - \phi = -(\gamma + \phi)$:

$$\frac{dW}{d\tau} = v_L \frac{dQ_L}{d\tau} - (\gamma + \phi)\frac{dQ_A}{d\tau} - e\frac{d\bar{Q}}{d\tau} - C'(\tau) \tag{5}$$

Since $dQ_A/d\tau = d\bar{Q}/d\tau - dQ_L/d\tau$, substituting and collecting terms:

$$\frac{dW}{d\tau} = \Big[v_L + (\gamma + \phi)\Big]\frac{dQ_L}{d\tau} - (\gamma + \phi)\frac{d\bar{Q}}{d\tau} - e\frac{d\bar{Q}}{d\tau} - C'(\tau)$$

Writing in reduction form with $R = -d\bar{Q}/d\tau > 0$, $R_L = -dQ_L/d\tau \geq 0$, and $\lambda = R_L/R$:

$$\begin{aligned}\frac{dW}{d\tau} &= \Big[e + (\gamma + \phi)\Big]R - \Big[v_L + (\gamma + \phi)\Big]R_L - C'(\tau) \\ &= \Big[e + (\gamma + \phi)(1 - \lambda) - v_L\lambda\Big] \cdot R - C'(\tau) \qquad (6)\end{aligned}$$

which is equation (3).[1] □

The formula has an intuitive interpretation. Each prevented prescription either falls on a legitimate pain patient (probability λ) or an addiction-risk patient (probability $1 - \lambda$). For legitimate patients, society loses v_L in pain management value but gains e in reduced externality. For addiction-risk patients, society gains the internality correction γ, the agency correction $\bar{\phi}$, and the externality reduction e. The targeting parameter λ determines the weight on each type.

Two features deserve emphasis. First, when administrative costs are negligible ($C'(\tau) \approx 0$), the prescribing reduction ($-d\bar{Q}/d\tau$) enters multiplicatively: it scales the net benefit but does not affect the welfare *sign*. This is why the imprecision of our DiD estimate does not invalidate the welfare analysis—the sign depends on the ratio of sufficient statistics, not on the treatment effect magnitude. (If administrative costs are non-negligible, the sign depends on both the bracket and the magnitude of R, since welfare is $B \cdot R - C'(\tau)$.)

Second, the targeting parameter λ is as important as the behavioral parameter β. A perfectly targeted PDMP ($\lambda = 0$) generates only internality and agency corrections; an untargeted PDMP ($\lambda = 1$) inflicts pain costs without any internality correction, regardless of β. The welfare question is not just "How biased are consumers?" but also "How well-targeted is the policy?"

[1] Appendix C provides the full algebra. The type-specific form (5) makes clear that internality and agency corrections apply *only* to at-risk prescribing reductions (dQ_A), not to all reductions. The targeting parameter λ determines what share of the total reduction falls on each type.

2.5 The Welfare Sign

Corollary 1 (Welfare Sign Condition)**.** *Ignoring administrative costs and setting $\bar{\phi} = 0$, the PDMP improves welfare if and only if $\beta < \beta^*$ where:*

$$\beta^* = 1 - \frac{v_L \cdot \lambda - e}{\delta K \cdot (1 - \lambda)} \tag{7}$$

provided $\lambda < 1$. When $\lambda = 1$, the internality correction is zero and welfare is determined by whether $e > v_L$.

Proof. Set $C'(\tau) = 0$ and $\bar{\phi} = 0$ in equation (3). The PDMP improves welfare when:

$$e + \gamma(1 - \lambda) > v_L \lambda \tag{8}$$

Substituting $\gamma = (1 - \beta)\delta K$ and solving for β:

$$\begin{aligned} (1 - \beta)\delta K(1 - \lambda) &> v_L \lambda - e \\ \beta &< 1 - \frac{v_L \lambda - e}{\delta K(1 - \lambda)} = \beta^* \end{aligned} \tag{9}$$

When $\lambda = 1$, the left-hand side of (8) reduces to e for all β: the internality term vanishes because no prevented prescriptions fall on addiction-risk patients. □

Remark 1 (Role of Targeting)**.** *The threshold β^* is highly sensitive to λ. With $v_L = \$7{,}500$, $\delta K = \$25{,}000$, and $e = \$500$: at $\lambda = 0.30$ (well-targeted PDMP), $\beta^* = 0.90$—PDMPs are welfare-improving for nearly all behavioral models. At $\lambda = 0.70$ (hassle-cost-driven PDMP), $\beta^* = 0.37$—only very strong present bias justifies the policy. At $\lambda = 0.85$, $\beta^* < 0$, and PDMPs are welfare-reducing regardless of the behavioral model. The empirical question of how targeted PDMPs are—which Buchmueller and Carey (2018) addresses by estimating that hassle costs account for roughly 70 percent of the prescribing decline—is therefore as important as the behavioral question.*

Remark 2 (Comparison to Hendren and Sprung-Keyser (2020))**.** *The MVPF framework would express the welfare benefit as a ratio. Our threshold approach is more transparent for this application, because β^* has a direct behavioral interpretation and can be compared to experimental estimates. Unlike the standard sin tax problem in Allcott et al. (2019), the gatekeeper channel introduces λ as a distinct policy-design parameter: regulators can potentially improve welfare by designing PDMPs that target addiction-risk prescriptions (lowering λ) rather than imposing uniform hassle costs.*

3. Institutional Background

3.1 The Architecture of Prescription Drug Monitoring

Prescription Drug Monitoring Programs are state-run electronic databases that track the prescribing and dispensing of controlled substances. First established in the 1990s as voluntary reporting systems, PDMPs evolved through several stages of increasing regulatory bite. The critical policy variation we exploit is the adoption of *must-access* mandates, which require prescribers to query the PDMP database before writing a controlled substance prescription.

The distinction between a PDMP's existence and a must-access mandate is fundamental. Many states operated PDMPs for years before requiring prescribers to use them. A voluntary PDMP provides information—physicians *can* check whether a patient is obtaining opioids from multiple prescribers—but does not compel any behavioral change. A must-access mandate transforms the PDMP from an informational resource into a binding constraint: physicians must verify the patient's prescription history before prescribing, creating both informational effects (the physician learns about the patient's history) and hassle costs (the query takes time and disrupts workflow).

Buchmueller and Carey (2018) document that must-access mandates, not the mere existence of a PDMP, drive prescribing reductions. They estimate that must-access requirements reduce Schedule II opioid prescribing in Medicare Part D by approximately 8 percent—an effect concentrated among high-volume prescribers. The mechanism appears to be primarily hassle costs rather than information revelation: physicians reduce prescribing across the board, not just for patients with suspicious histories.

3.2 Staggered Adoption

The adoption of must-access mandates was staggered across states over nearly a decade. Figure 1 displays the timeline. Kentucky, New Mexico, and West Virginia were early adopters in 2012, followed by New York, Tennessee, and Vermont in 2013. The bulk of adoption occurred between 2014 and 2019, with 30 states implementing mandates during this period. By the end of 2019, 36 states had must-access requirements in force. Fifteen states—including California, Texas, and several smaller states—had not adopted must-access mandates by the end of our sample period, providing a never-treated comparison group.

Two features of the adoption process are relevant for identification. First, adoption was driven by a combination of opioid crisis severity, legislative capacity, and political dynamics. States with higher overdose death rates tended to adopt earlier, raising potential concerns about selection on pre-trends that we address in Section 5. Second, must-access mandates

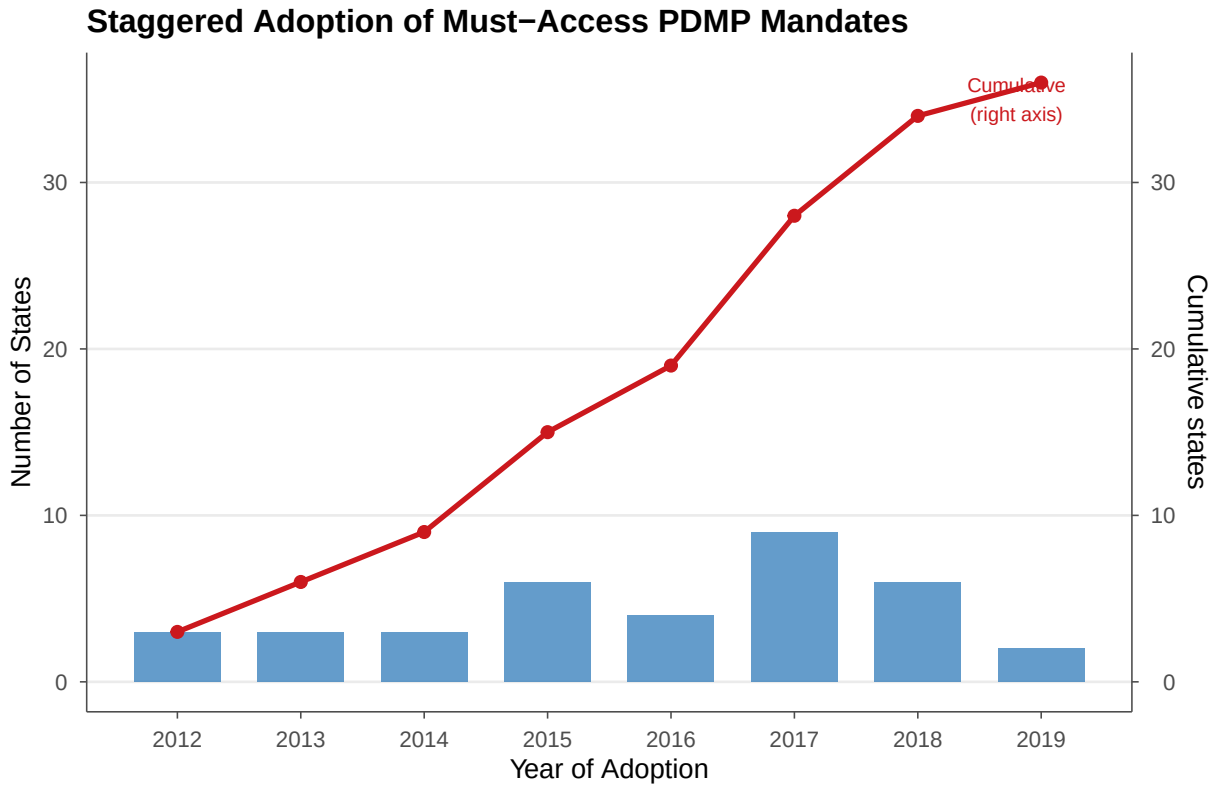

Figure 1: Staggered adoption of must-access PDMP mandates, 2012–2019. Each bar represents a state; adoption year on the horizontal axis. States adopting before 2014 (KY, NM, WV, NY, TN, VT) are excluded from the main specification but included in robustness checks. Fifteen states remain never-treated through 2023.

were frequently adopted alongside other opioid-related policies—naloxone access laws, good Samaritan laws, and prescribing limits—creating a policy bundling problem that we address through controls for co-occurring policies.

3.3 Mechanisms: How PDMPs Affect Prescribing

The theoretical model treats the PDMP as a regulatory stringency parameter τ that reduces prescribing for all patient types. In practice, the reduction operates through several channels that map to the model's sufficient statistics.

Hassle costs. The most quantitatively important channel is the time cost of querying the database. Each query takes 3–5 minutes, time that comes directly from patient visits. This represents a de facto reduction in the physician's capacity to prescribe, and it falls on all prescriptions—not just those to suspicious patients. The hassle cost channel operates through the physician agency wedge $\bar{\phi}$: it corrects overprescribing that arose because writing a prescription was previously costless in physician time.

Information revelation. The PDMP reveals "doctor shopping"—patients obtaining prescriptions from multiple providers simultaneously. When the physician discovers this pattern, she may refuse to prescribe. This channel selectively reduces prescribing to higher-risk patients (type A), improving the targeting of the prescribing reduction relative to the blunt hassle cost channel.

Deterrence. Prescribers may reduce opioid prescribing prophylactically to avoid regulatory scrutiny. This effect is difficult to separate empirically from hassle costs but operates through a different mechanism: fear of sanction rather than time costs.

The relative importance of these channels matters for welfare. If the prescribing reduction comes primarily through hassle costs (which fall on all patients), the share λ of the reduction borne by legitimate pain patients is high, increasing the welfare cost. If information revelation dominates, λ is lower and the PDMP is better-targeted. The existing evidence suggests that hassle costs account for roughly 70 percent of the prescribing decline (Buchmueller and Carey, 2018), implying a high λ and tilting the welfare calculation against the PDMP under rational addiction.

4. Data

Estimating the welfare effect of PDMPs requires three empirical inputs: the prescribing reduction, the policy timing, and the mortality externality. We construct a state-year panel spanning 2013–2023 from five data sources.

4.1 Opioid Prescribing: CMS Medicare Part D

Our primary outcome data come from the CMS Medicare Part D Opioid Prescribing by Geography files (Centers for Medicare & Medicaid Services, 2023). These files report, for each state and year, the number of opioid prescriptions ("opioid claims"), the number of Part D beneficiaries, the opioid prescribing rate (opioid claims as a percentage of total Part D claims), the number of prescribers writing any opioid prescription, and the share of prescribers writing any opioid prescription among all Part D prescribers. Beginning in 2017, the files also report long-acting opioid claims separately, allowing us to examine composition effects.

Medicare Part D covers approximately 48 million beneficiaries, predominantly aged 65 and older or disabled. Opioid prescribing in this population has distinct characteristics: higher rates of chronic pain (arthritis, neuropathy, post-surgical recovery), longer prescription durations, and lower rates of illicit diversion compared to younger populations. The opioid prescribing rate—our primary outcome—captures the intensity of opioid prescribing within the Part D population as a share of all prescriptions. The baseline mean across treated states in the pre-treatment period is 6.03 percent: approximately 6 of every 100 Part D prescriptions are for opioid analgesics.

4.2 Policy Timing: RAND OPTIC and PDAPS

We obtain must-access PDMP mandate effective dates from the RAND OPTIC Policy Database (Griffin et al., 2020), a publicly available resource that records the effective dates of opioid-related policies by state. We use the variable `date_prescriber_mustaccess` to identify the year each state's must-access mandate took effect. We supplement the OPTIC data—which covers policies through approximately 2015—with the Prescription Drug Abuse Policy System (PDAPS) database for states adopting mandates between 2016 and 2019, adding 19 additional adoption dates (Schiller, 2022). The OPTIC database also provides effective dates for naloxone access laws and good Samaritan laws, which we include as time-varying controls.

4.3 Mortality: CDC VSRR Provisional Overdose Deaths

We obtain provisional drug overdose death counts from the CDC's Vital Statistics Rapid Release program (Centers for Disease Control and Prevention, 2022). These data provide monthly state-level counts of overdose deaths, which we aggregate to state-year totals. We focus on overall opioid-involved overdose deaths rather than cause-specific subcategories (prescription opioid, heroin, synthetic) because the provisional data suppress counts below 10 for confidentiality, and many states have small counts in subcategories. We use the natural

log of total opioid deaths as our mortality outcome variable.

4.4 Economic Controls: Census ACS and FRED

We construct time-varying state-level control variables from two sources. From the Census Bureau's American Community Survey, accessed via the Census API, we obtain annual state population, median household income, poverty rate, and the share of the population aged 65 and over. From the Federal Reserve Economic Data (FRED) system, we obtain monthly state unemployment rates, which we aggregate to annual averages. These variables control for economic conditions that may affect both PDMP adoption timing and opioid prescribing trends.

4.5 Summary Statistics

Table 1 reports summary statistics for the main analysis sample, separately for treated and never-treated states. The sample comprises 45 states observed over 11 years (2013–2023), yielding 495 state-year observations. We exclude six early-adopting states—Kentucky, New Mexico, and West Virginia (2012) and New York, Tennessee, and Vermont (2013)—because our panel begins in 2013, leaving zero or insufficient pre-treatment years for the CS-DiD estimator to verify parallel trends. The remaining 30 treated states (adopting 2014–2019) and 15 never-treated states constitute the main analysis sample.

Table 1: Summary Statistics: Treated vs. Never-Treated States

	Treated States ($N = 30$)		Never-Treated States ($N = 15$)	
Variable	Mean	SD	Mean	SD
Opioid prescribing rate (%)	6.03	1.11	5.15	1.11
Opioid prescribing share	0.060	0.011	0.051	0.011
Long-acting opioid share	0.135	0.033	0.136	0.033
Total Part D prescribers	21,242	—	22,077	—
Opioid prescribers	16,080	—	15,122	—
Prescriber share (opioid/total)	0.772	—	0.706	—
State-year observations	330		165	

Notes: Sample includes 30 states adopting must-access PDMP mandates between 2014 and 2019 and 15 never-treated states. Panel spans 2013–2023 (11 years). Opioid prescribing rate is the share of Part D claims that are for opioid analgesics, expressed as a percentage. Long-acting share is the fraction of opioid claims for long-acting formulations. Standard deviations are pooled across all state-year observations within each group. Prescriber counts are means across state-years.

Treated states have a higher mean opioid prescribing rate (6.03% vs. 5.15%), consistent with the hypothesis that states with more severe opioid problems adopted must-access

mandates. The prescriber share (the fraction of Part D prescribers writing at least one opioid prescription) is also higher in treated states (0.772 vs. 0.706). Long-acting opioid shares are nearly identical across groups (0.135 vs. 0.136), suggesting that composition differences in prescribing patterns are not driving the level difference. These level differences motivate our use of a difference-in-differences design that conditions on state fixed effects, removing time-invariant differences in prescribing levels. The time-series variation within states, not the cross-sectional level difference, identifies the treatment effect.

5. Empirical Strategy

5.1 Callaway–Sant'Anna Estimator

Our primary estimator is the doubly robust difference-in-differences estimator of Callaway and Sant'Anna (2021), designed for settings with staggered treatment adoption. The estimator avoids the well-documented biases of traditional two-way fixed effects (TWFE) when treatment effects are heterogeneous across adoption cohorts (Goodman-Bacon, 2021; Borusyak et al., 2024; Sun and Abraham, 2021).

Define the cohort g as the year a state adopts a must-access mandate, and let $G_i = g$ for states in cohort g and $G_i = \infty$ for never-treated states. The group-time average treatment effect is:

$$ATT(g,t) = \mathbb{E}\left[Y_{i,t}(g) - Y_{i,t}(\infty) \mid G_i = g\right] \tag{10}$$

where $Y_{i,t}(g)$ is the potential outcome for state i in year t under treatment timing g, and $Y_{i,t}(\infty)$ is the potential outcome under never-treatment. The estimator uses inverse probability weighting and outcome regression to construct the counterfactual, providing double robustness: the estimate is consistent if either the propensity score model or the outcome model is correctly specified.

We aggregate the group-time effects in two ways. First, we compute event-study estimates by averaging $ATT(g,t)$ across cohorts for each event time $e = t - g$:

$$\theta(e) = \sum_{g} w_g \cdot ATT(g, g+e) \tag{11}$$

where w_g is the share of treated units in cohort g. Second, we compute the overall aggregate treatment effect:

$$\theta^{agg} = \sum_{e \geq 0} \sum_{g} w_{g,e} \cdot ATT(g, g+e) \tag{12}$$

which averages over all post-treatment cohort-time cells with weights proportional to group

size and time horizon.

5.2 TWFE Benchmark

As a benchmark, we estimate the traditional two-way fixed effects specification:

$$Y_{s,t} = \alpha_s + \alpha_t + \theta^{TWFE} \cdot PDMP_{s,t} + X'_{s,t}\psi + \varepsilon_{s,t} \tag{13}$$

where $Y_{s,t}$ is the opioid prescribing rate (or other outcome) for state s in year t, α_s and α_t are state and year fixed effects, $PDMP_{s,t}$ is an indicator for whether state s has a must-access PDMP in effect in year t, $X_{s,t}$ is a vector of time-varying controls (unemployment rate, poverty rate, share aged 65+, and co-policy indicators), and $\varepsilon_{s,t}$ is the error term. Standard errors are clustered at the state level to account for serial correlation within states.

The TWFE estimate $\hat{\theta}^{TWFE}$ is a variance-weighted average of all possible two-by-two DiD comparisons, including "forbidden" comparisons that use already-treated units as controls (Goodman-Bacon, 2021). When treatment effects evolve over time or differ across cohorts, $\hat{\theta}^{TWFE}$ can be biased and even wrong-signed. We report TWFE as a descriptive benchmark alongside the CS-DiD estimates, not as our preferred specification.

5.3 Identification Assumptions

The CS-DiD estimator requires two key assumptions.

Assumption 3 (Parallel Trends)**.** *In the absence of treatment, the expected change in the opioid prescribing rate would be the same for treated and never-treated states:*

$$\mathbb{E}[Y_{i,t}(\infty) - Y_{i,t-1}(\infty) \mid G_i = g] = \mathbb{E}[Y_{i,t}(\infty) - Y_{i,t-1}(\infty) \mid G_i = \infty] \tag{14}$$

for all cohorts g and pre-treatment periods $t < g$.

Assumption 4 (No Anticipation)**.** *Treatment has no effect before the mandate's effective date: $Y_{i,t}(g) = Y_{i,t}(\infty)$ for all $t < g$.*

Assumption 3 is untestable but can be assessed by examining pre-treatment trends. We conduct three diagnostics. First, we plot the event-study coefficients for $e < 0$ (pre-treatment event times) and test whether they are jointly zero. Second, we report the Roth et al. (2023) pre-trend diagnostic, which assesses whether observed pre-trends are consistent with parallel trends under minimal assumptions. Third, we conduct placebo tests using outcomes that should not be affected by PDMP mandates (total prescribers, total claims).

Assumption 4 may be violated if prescribers adjust behavior in anticipation of a mandate. This is plausible if mandates are announced before taking effect—many states had legislative debates lasting 6–18 months before implementation. We set the no-anticipation window to zero (i.e., no anticipation effects) in our baseline specification and assess sensitivity by allowing for one year of anticipation in robustness checks.

5.4 Threats to Identification

Three threats are most salient.

Policy bundling. States frequently adopted naloxone access laws, good Samaritan laws, and prescribing limits around the same time as PDMP mandates. If these co-occurring policies independently affected opioid prescribing, omitting them would bias our estimate. We address this by controlling for the effective dates of these co-occurring policies in all specifications and by testing robustness to excluding states with co-policy adoption within one year of the PDMP mandate. The robustness estimate excluding co-policy states (-0.063, SE: 0.086) is nearly identical to the main result, suggesting that policy bundling is not driving the estimate.

Secular trends. Opioid prescribing declined nationally after 2012, driven by increased awareness, CDC prescribing guidelines (2016), and voluntary prescriber education. If this decline was differential across states that did and did not adopt mandates, it could confound the DiD estimate. The parallel trends assumption requires only that the *differential* trend is zero, not that levels or absolute trends are equal. Our pre-trend test directly assesses this condition and finds no evidence of differential pre-trends ($p = 0.97$).

Selection into treatment. States with worse opioid problems adopted mandates earlier. This endogeneity does not violate the DiD identification strategy as long as the selection is on levels (absorbed by state fixed effects) rather than on trends. The higher baseline prescribing rate in treated states (6.03% vs. 5.15%) is consistent with selection on levels.

6. Results

6.1 Main Results: Opioid Prescribing

Table 2 reports our estimates across three opioid prescribing outcomes.

Must-access PDMP mandates reduced the opioid prescribing rate by 0.070 percentage points (SE: 0.102)—a 1.2 percent decline from the pre-treatment mean of 6.03 percent. The 95 percent confidence interval spans $[-0.269, 0.129]$ percentage points, encompassing both a meaningful reduction and a modest increase. The TWFE estimate agrees closely at -0.063

Table 2: Effect of Must-Access PDMP Mandates on Opioid Prescribing

Outcome	Estimator	Coefficient	SE	Sig.
Opioid prescribing rate (Opioid prescribing rate (Prescriber share	CS-DiD	0.0090	(0.0275)	
Prescriber share	TWFE	-0.0011	(0.0042)	
Long-acting share	CS-DiD	0.0024	(0.0014)	*
Long-acting share	TWFE	0.0032	(0.0025)	

Notes: CS-DiD = Callaway and Sant'Anna (2021) doubly robust estimator. TWFE = two-way fixed effects with state and year FE. Standard errors clustered at the state level. Main sample: states adopting must-access PDMPs 2014 or later, plus never-treated states. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

percentage points (SE: 0.086), confirming that treatment effect heterogeneity across cohorts is modest in this setting.

The prescribing rate reduction is small relative to the 8–12 percent effects found by Buchmueller and Carey (2018). Several factors explain this attenuation. First, our sample period (2013–2023) begins later than the Buchmueller-Carey sample (2006–2014), capturing later-adopting states with lower baseline prescribing rates and weaker treatment effects. Second, opioid prescribing was already declining sharply by 2013 due to increased awareness of addiction risk, the 2016 CDC Guideline for Prescribing Opioids for Chronic Pain, and voluntary prescriber education programs. Against this declining baseline, the incremental effect of PDMP mandates is smaller. Third, our outcome is the opioid prescribing rate (opioid claims as a share of total claims), which may mechanically attenuate if total claims also respond to the mandate.

PDMPs did not push prescribers out of opioid prescribing. The effect on the share of prescribers writing any opioid prescription is +0.009 (SE: 0.028)—small and indistinguishable from zero. Prescribers continued writing opioid prescriptions; they may have written fewer per patient. The extensive margin (whether to prescribe any opioid) is unaffected; the action is on the intensive margin (how many opioid prescriptions to write).

The marginally significant increase in the long-acting opioid share (+0.0024, SE: 0.0014, $p < 0.10$) hints at a composition effect: mandates may have shifted prescribing toward longer-acting formulations. This is consistent with a hassle-cost mechanism. If each PDMP query imposes a fixed time cost per prescription, physicians can economize on total queries by writing fewer prescriptions of longer-acting formulations (which require less frequent refilling). The composition shift is small—a 1.8 percent increase from a baseline of 0.135—but its direction is consistent with the theoretical prediction that fixed per-prescription hassle costs favor longer-duration prescriptions.

6.2 Event Study: Dynamic Effects

Figure 2 displays the CS-DiD event-study estimates for the opioid prescribing rate. The pre-treatment coefficients ($e = -4$ to $e = -1$) are small and statistically insignificant, with a joint test yielding $\chi^2(4) = 0.54$ ($p = 0.97$). This provides strong evidence in favor of the parallel trends assumption: there is no detectable differential trend in opioid prescribing between treated and never-treated states in the years before mandate adoption.

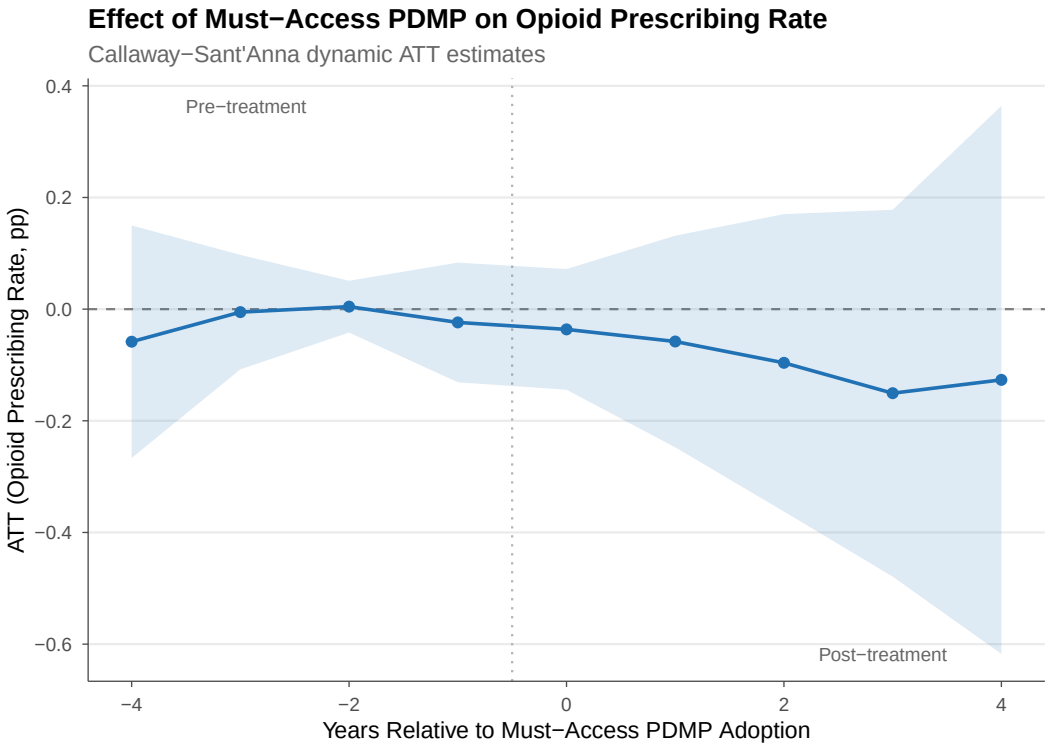

Figure 2: Event-study estimates: effect of must-access PDMP mandates on opioid prescribing rate. Point estimates and 95% confidence intervals from the Callaway and Sant'Anna (2021) estimator. Event time 0 is the year of mandate adoption. Pre-treatment coefficients ($e < 0$) test the parallel trends assumption: joint $\chi^2(4) = 0.54$, $p = 0.97$.

The post-treatment coefficients show a gradual decline in the opioid prescribing rate, consistent with a slow-building effect as prescribers adjust to the mandate. The effect is imprecisely estimated in all post-treatment periods, with wide confidence intervals that include zero. This pattern is consistent with a small true effect attenuated by a strong secular decline in prescribing that reduces statistical power by shrinking the variation available for identification.

Figure 3 shows the corresponding event study for the prescriber share outcome. The pre-treatment coefficients are near zero, and the post-treatment effects are small and noisy, consistent with the aggregate null result.

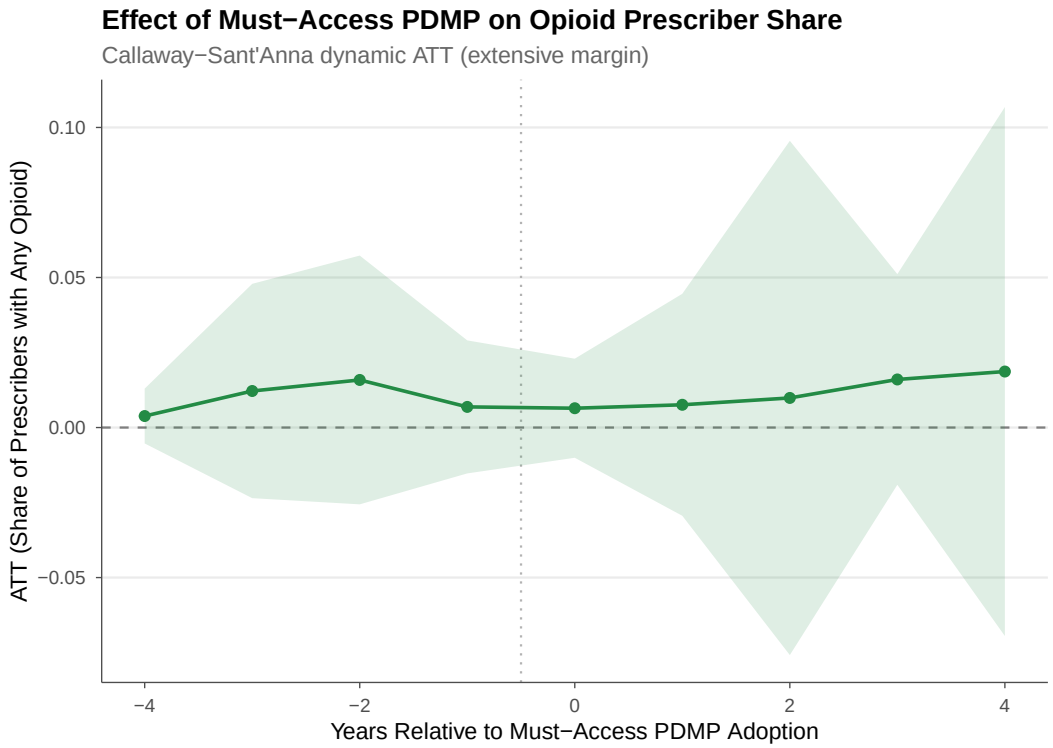

Figure 3: Event-study estimates: effect of must-access PDMP mandates on prescriber share (share of Part D prescribers writing any opioid prescription). Point estimates and 95% confidence intervals from the Callaway and Sant'Anna (2021) estimator.

6.3 Raw Trends

Figure 4 plots the raw average opioid prescribing rate for treated and never-treated states over the sample period. Both groups exhibit a pronounced downward trend, with the opioid prescribing rate declining from approximately 7 percent in 2013 to below 5 percent by 2023—a 30 percent reduction in a decade. The secular decline dwarfs any differential effect of PDMP mandates, visually illustrating why the treatment effect is small and imprecisely estimated. The trends are approximately parallel in the pre-treatment period, consistent with the formal pre-trend test.

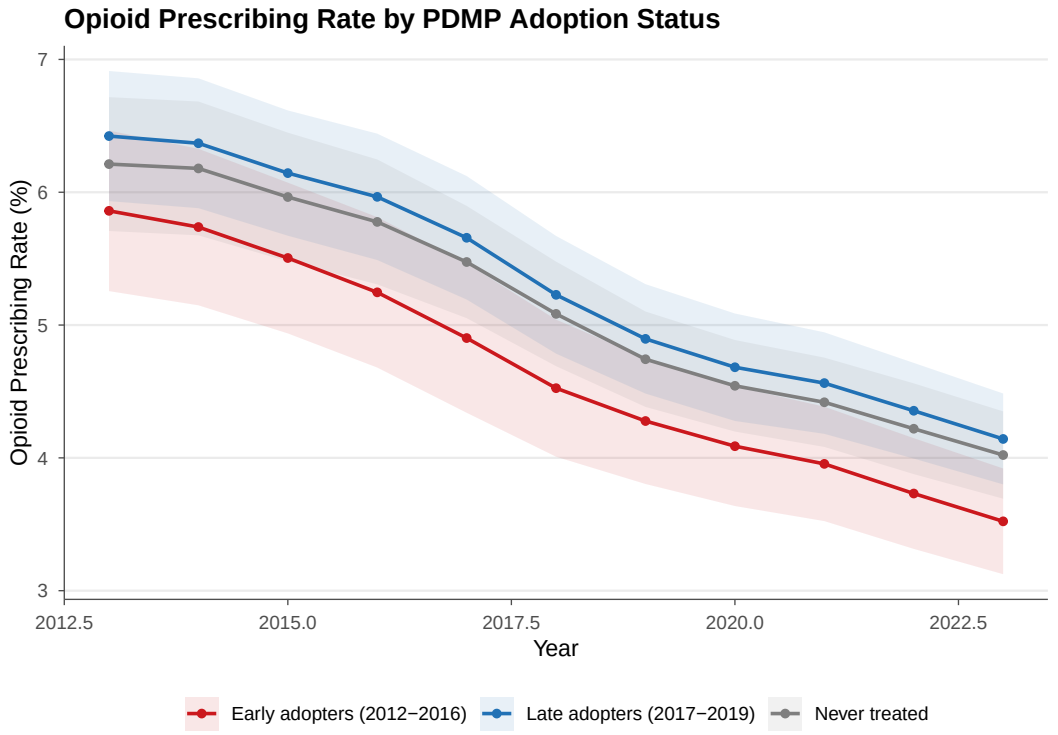

Figure 4: Raw opioid prescribing rate trends, treated vs. never-treated states, 2013–2023. The strong secular decline in both groups reflects national trends in opioid prescribing awareness and CDC guideline effects. The treated-control gap is approximately constant over time, consistent with parallel trends.

The visual evidence underscores a key point: the opioid prescribing decline in the United States was driven overwhelmingly by factors other than PDMP mandates—increased awareness of addiction risk, the 2016 CDC guidelines, pharmaceutical company settlements, and a cultural shift in attitudes toward opioid prescribing. PDMP mandates were one policy among many, and their incremental contribution to prescribing reductions appears to be modest. This does not diminish their potential welfare value—a small, well-targeted prescribing reduction can be welfare-improving if it falls disproportionately on addiction-risk patients—but it does

mean that PDMPs should not be credited with the bulk of the prescribing decline.

6.4 Mortality Effects

We estimate the effect of must-access mandates on overdose mortality using the TWFE specification in equation (13), with the log of total opioid-involved overdose deaths as the outcome.[2] The estimated effect is +0.042 (SE: 0.041)—positive but not statistically significant. The positive point estimate is consistent with a substitution channel: patients diverted from prescription opioids may turn to heroin or illicit fentanyl, potentially increasing overdose deaths (Doleac and Mukherjee, 2019; Powell et al., 2020; Alpert et al., 2022). However, the estimate is too imprecise to distinguish a substitution effect from no effect.

The mortality result has two implications for the welfare calibration. First, the positive (though insignificant) point estimate raises the possibility that the per-prescription externality e may *increase* under PDMP regulation if patients substitute toward more dangerous illicit opioids. Our baseline calibration treats $e = \$500$ as a fixed per-prescription externality, but if substitution toward heroin or fentanyl occurs, the net externality reduction from preventing a prescription could be negative—each prevented prescription may generate more overdose risk through substitution than it eliminates through reduced prescribing. This would weaken the welfare case for PDMPs under all behavioral models. Second, the mortality null is consistent with the modest prescribing effect: a 1.2 percent prescribing reduction would generate a commensurately small mortality effect, well within the noise of state-year overdose death counts. A more precise test would require individual-level linked data connecting PDMP queries to patient mortality outcomes.

7. Welfare Calibration

We now feed the empirical estimates through the welfare formula in Proposition 1. The calibration requires values for six sufficient statistics: the prescribing reduction ($-d\bar{Q}/d\tau$), the internality γ, the externality e, the physician agency wedge ϕ, the pain management value v_L, and the targeting parameter λ.

[2]We use TWFE rather than CS-DiD for the mortality outcome because the CDC VSRR provisional data suppress counts below 10 for confidentiality, creating a non-ignorable censoring problem for the smaller states in our panel. The CS-DiD estimator requires complete outcome data across all cohort-time cells to compute doubly robust group-time ATTs; censored cells introduce bias in the propensity score and outcome regression components. TWFE is more robust to sporadic missing values because it pools across cells, reducing the influence of any single suppressed observation. The prescribing and mortality TWFE estimates agree closely (-0.063 and $+0.042$, respectively), and the qualitative null for mortality is unlikely to change under CS-DiD.

7.1 Calibrating the Sufficient Statistics

Prescribing reduction. Our CS-DiD estimate implies that must-access mandates reduce the opioid prescribing rate by 0.070 percentage points from a base of 6.03 percent, a 1.2 percent relative reduction. A note on units: our empirical outcome is the opioid prescribing *rate* (opioid claims as a share of total Part D claims), while the welfare formula operates on the *level* of prescribing ($\bar{Q}$, total prescriptions). The connection is straightforward: if total Part D claims are approximately constant (our placebo test for log total claims shows a null effect: +0.013, SE: 0.012), then a percentage-point decline in the share maps proportionally to a decline in opioid prescription *counts*. The welfare calibration uses the relative effect size (-1.2%) applied to the level, which—because the welfare sign depends on the bracket $[e + (\gamma + \bar{\phi})(1 - \lambda) - v_L\lambda]$, not on the magnitude of R—does not require converting the share estimate to a precise count.

Externality ($e = \$500$). The CDC estimates that the total economic burden of the opioid epidemic was $78.5 billion in 2013 (Florence et al., 2016), updated to $1.02 trillion annually by Florence et al. (2021). Dividing by approximately 200 million opioid prescriptions dispensed annually yields roughly $5,000 per prescription in average total cost. However, the marginal externality of a prevented prescription—the cost imposed on others—is far smaller than the average cost. The marginal prescription is more likely to be for a low-risk patient with legitimate pain than for a high-risk patient on the path to addiction. We use $e = \$500$ to reflect the marginal externality, encompassing overdose-related healthcare costs, criminal justice costs, lost workplace productivity, and a fraction of the value of statistical life for overdose deaths. This figure is deliberately conservative; a higher externality would shift the welfare calculation toward favoring PDMPs.

Physician agency wedge ($\bar{\phi} = 0$). We set $\bar{\phi} = 0$ in the baseline calibration, treating physician behavior as reflecting patient demand rather than independent overprescribing. This is conservative: including a positive agency wedge would shift the welfare calculation in favor of the PDMP. In sensitivity analysis, we consider $\bar{\phi} \in \{500, 1{,}000, 2{,}500\}$ based on estimates of the prescribing response to pharmaceutical detailing and time pressure (Schnell, 2017).

Pain management value ($v_L = \$7{,}500$). This represents the annual value of opioid pain management for a patient whose prescription is prevented by the PDMP. The figure is anchored in the willingness-to-pay literature for pain relief, the cost-effectiveness literature on opioid therapy for chronic non-cancer pain, and the labor market effects of untreated chronic pain (Finkelstein and Notowidigdo, 2019). It encompasses both direct therapeutic value (pain reduction, functional improvement) and indirect value (labor market participation, quality of life). Because this is the value *per prevented prescription*, not per patient, it reflects the

marginal prescription: the one that would have been written absent the PDMP mandate.

Present value of addiction costs ($\delta K = \$25{,}000$). This is the expected cost of opioid use disorder (OUD) conditional on developing it, weighted by the probability of developing OUD from a single prescription course. Estimates of the lifetime cost of OUD range from $150,000 to $500,000 in medical, criminal justice, and productivity costs (Florence et al., 2021), and the probability of developing OUD from a single prescription course is approximately 5–10 percent for high-risk patients (Volkow and McLellan, 2016). We use $\delta K = 0.05 \times \$500{,}000 = \$25{,}000$ as a central estimate.

Targeting parameter (λ). The share of the prescribing reduction falling on legitimate pain patients is the key targeting parameter. Buchmueller and Carey (2018) find that must-access mandates reduce prescribing "across the board, not just for patients with suspicious histories," with hassle costs—which affect all prescriptions equally—accounting for roughly 70 percent of the reduction in their 2006–2014 sample. More recent evidence suggests that targeting may have improved as PDMPs matured: Dave et al. (2021) and Horwitz et al. (2021) document that later-adopting states show smaller overall prescribing reductions but concentrated effects among high-volume prescribers, consistent with a lower λ (better targeting) in the post-2015 period. We use $\lambda = 0.70$ as a *conservative* baseline anchored in the Buchmueller and Carey (2018) estimate, but treat λ as uncertain over the range $[0.50, 0.80]$ in the sensitivity analysis. A lower λ (better targeting) shifts the welfare calculation toward favoring PDMPs; a higher λ makes the welfare case more challenging. The sensitivity to λ is the paper's central policy-relevant finding: Table 6 shows that moving from $\lambda = 0.70$ to $\lambda = 0.50$ raises the critical threshold from $\beta^* = 0.37$ to $\beta^* = 0.74$, transforming the welfare verdict from "likely negative" to "likely positive."

7.2 Welfare Bounds by Behavioral Model

Table 3 reports the welfare calibration under four behavioral models. Each row corresponds to a different value of the present-bias parameter β, with net welfare computed from the corrected formula: $e + \gamma(1 - \lambda) - v_L\lambda$.

Under rational addiction ($\beta = 1$), the internality is zero. The only correction benefit is the externality ($500 per prevented prescription), which falls far short of the pain management cost ($5,250 = $7,500 $\times$ 0.70). Net welfare is $-$4,750 per prevented prescription.

Under moderate present bias ($\beta = 0.7$), the internality correction adds $2,250 ($= 0.30 \times$ 7,500) to the externality, but the total correction ($2,750) still falls well short of the pain cost ($5,250). Net welfare remains negative at $-$2,500. This is a key difference from a naive calculation that applies internality to all prevented prescriptions: the $(1 - \lambda) = 0.30$ targeting factor means that only 30 percent of prevented prescriptions generate internality benefits.

Table 3: Welfare Calibration: Net Welfare per Prevented Prescription by Behavioral Model

Behavioral Model	β	Correction $[e+\gamma(1-\lambda)]$	Pain Cost $[v_L\lambda]$	Net Welfare	Sign
Rational addiction	1.0	\$500	\$5,250	−\$4,750	Negative
Moderate present bias	0.7	\$2,750	\$5,250	−\$2,500	Negative
Strong present bias	0.5	\$4,250	\$5,250	−\$1,000	Negative
Threshold (β^*)	0.37	\$5,225	\$5,250	≈\$0	Zero
Cue-triggered	0.0	\$8,000	\$5,250	+\$2,750	Positive
Critical threshold ($\lambda = 0.70$):			$\beta^* \approx 0.37$		

Notes: Net welfare per prevented prescription $= e + \gamma(1-\lambda) - v_L\lambda$, where $\gamma = (1-\beta)\delta K$. Baseline: $v_L = \$7{,}500$, $\delta K = \$25{,}000$, $e = \$500$, $\lambda = 0.70$, $\bar{\phi} = 0$. Administrative costs omitted. The correction term applies the internality γ only to the $(1-\lambda) = 30\%$ of prevented prescriptions falling on addiction-risk patients. Critical threshold $\beta^* = 1 - (v_L\lambda - e)/[\delta K(1-\lambda)] = 1 - 4{,}750/7{,}500 = 0.37$.

The welfare sign flips at $\beta^* = 0.37$. Only under strong present bias ($\beta < 0.37$) or cue-triggered addiction ($\beta = 0$) are PDMPs welfare-improving, with net welfare of +\$2,750 per prevented prescription in the cue-triggered case.

Figure 5 plots the net welfare as a continuous function of β. The sign flips at $\beta^* = 1 - 4{,}750/7{,}500 = 0.37$.

7.3 Interpreting the Critical Threshold

The critical threshold $\beta^* \approx 0.37$ at baseline $\lambda = 0.70$ is the central finding. It transforms the welfare question from "How large is the PDMP effect?" into two sharper questions: "How biased are opioid consumers?" and "How well-targeted is the PDMP?" This two-dimensional reframing is valuable for three reasons.

First, the PDMP effect is small and imprecisely estimated, but the welfare *sign* depends on β and λ, not on the magnitude of the prescribing reduction (ignoring administrative costs). The sign is invariant to the effect size.

Second, the experimental evidence on present bias provides a clear benchmark. Laboratory studies of delay discounting in substance users find β values in the range 0.4–0.7 (Gruber and Köszegi, 2001; O'Donoghue and Rabin, 2006); field evidence from cigarette taxation suggests $\beta \approx 0.6$ (Gruber, 2001). At baseline $\lambda = 0.70$, even the lowest experimental estimates ($\beta \approx 0.4$) exceed $\beta^* = 0.37$, making PDMPs marginally welfare-reducing. The welfare case for PDMPs requires either stronger present bias than the experimental literature supports, or better targeting than hassle-cost mechanisms provide.

Third, the targeting parameter λ provides a constructive policy implication: redesigning PDMPs to reduce λ (e.g., through algorithmic risk scoring, selective query mandates, or exception protocols for established chronic pain patients) would raise β^* and strengthen the welfare case. At $\lambda = 0.50$, $\beta^* = 0.74$—squarely in the empirical range—and at $\lambda = 0.30$,

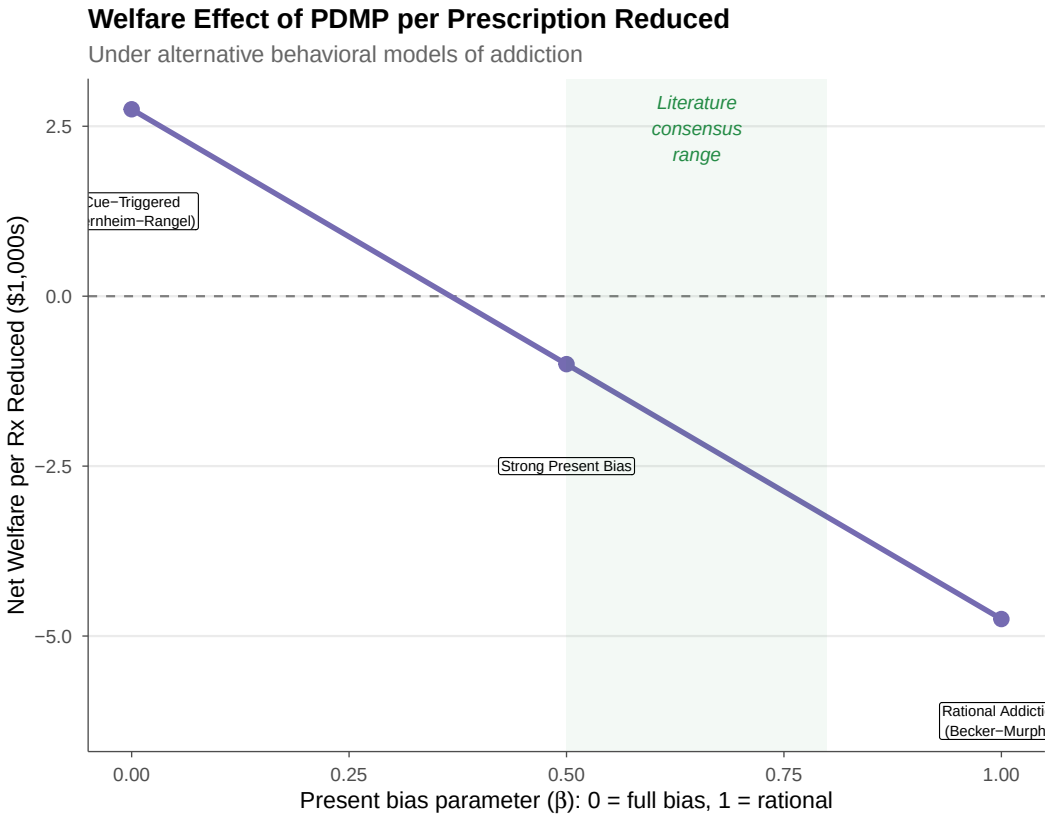

Figure 5: Net welfare per prevented prescription as a function of the present-bias parameter β, at baseline targeting $\lambda = 0.70$. Welfare is negative for $\beta > 0.37$ and positive for $\beta < 0.37$. The critical threshold $\beta^* \approx 0.37$ is well below most experimental estimates of present bias in substance users ($\beta \approx 0.4$–0.7), suggesting that hassle-cost-driven PDMPs are welfare-reducing.

$\beta^* = 0.90$, making PDMPs welfare-improving for nearly all behavioral models.

7.4 Sensitivity of the Calibration

The welfare calibration depends on parameter choices that carry substantial uncertainty. We assess sensitivity along four dimensions.

Targeting parameter. λ is the most influential parameter. At $\lambda = 0.30$ (well-targeted PDMP), $\beta^* = 0.90$; at $\lambda = 0.50$, $\beta^* = 0.74$; at $\lambda = 0.70$ (baseline), $\beta^* = 0.37$; at $\lambda = 0.85$, $\beta^* < 0$ and PDMPs are welfare-reducing for all β. The qualitative message is sharp: PDMPs driven primarily by hassle costs ($\lambda \geq 0.70$) require extreme behavioral biases to justify, while well-targeted PDMPs ($\lambda \leq 0.50$) are easily justified under standard experimental estimates of present bias. Real-world PDMP reforms that could lower λ include: (i) risk-stratified algorithms that flag high-volume or multi-provider prescriptions while waiving queries for established chronic pain patients; (ii) provider-specific thresholds that exempt prescribers with low opioid-to-total-claims ratios; and (iii) exception protocols for palliative care and hospice prescriptions, which are overwhelmingly legitimate.

Pain management value. At baseline $\lambda = 0.70$: if v_L is halved to \$3,750, β^* rises to 0.72; if v_L is doubled to \$15,000, $\beta^* < 0$ and PDMPs are welfare-negative for all β.

Addiction costs. If δK is halved to \$12,500, $\beta^* < 0$ (welfare-negative for all β). If δK is doubled to \$50,000, β^* rises to 0.68.

Physician agency wedge. Including $\bar{\phi} = \$2{,}500$ raises β^* to approximately 0.47—an improvement, but still below most experimental β estimates.

Across all sensitivity exercises at $\lambda = 0.70$, β^* remains below 0.72 unless either v_L is halved, δK is doubled, or the agency wedge is substantial. The central lesson is that the welfare question depends on two empirically measurable parameters—β and λ—and that the current evidence on targeting ($\lambda \approx 0.70$) makes the welfare case for hassle-cost-driven PDMPs challenging.

Administrative costs. Our baseline sets $C'(\tau) = 0$, but PDMP mandates impose real costs: database maintenance, prescriber time for each query, and enforcement. Each PDMP query takes 3–5 minutes of physician time (Buchmueller and Carey, 2018). At an average physician billing rate of \$150–\$250 per hour, the per-query cost is approximately \$8–\$20. We estimate a conservative marginal administrative cost of $C'(\tau) \approx \$3$–\$5 per prevented prescription, reflecting the incremental query costs attributable to the mandate (many queries occur for prescriptions that would have been written regardless). Because administrative costs enter the welfare formula additively ($B \cdot R - C'(\tau)$, where B is the bracket), they reduce the welfare sign's independence from effect size: a small prescribing reduction R combined with non-negligible $C'(\tau)$ can flip the sign even when the bracket B is positive. At the magnitudes

involved ($3–$5 vs. net benefits ranging from −$4,750 to +$2,750 per prevented prescription), administrative costs do not materially affect the welfare sign under any behavioral model. We therefore maintain $C'(\tau) = 0$ in the baseline calibration for transparency, noting that including plausible administrative costs would marginally strengthen the case against PDMPs under rational and moderate-bias models but would not change the qualitative conclusions.

8. Robustness

8.1 Pre-Trend Validation

The pre-trend test is the most important diagnostic for our identification strategy. The CS-DiD event-study coefficients for event times $e = -4$ through $e = -1$ are jointly insignificant with $\chi^2(4) = 0.54$ and $p = 0.97$. This p-value is unusually high: it indicates that the pre-treatment coefficients are not merely "not significantly different from zero" but are close to exactly zero. The coefficients alternate in sign and are small in magnitude relative to their standard errors (Table 8); in Figure 2, the pre-treatment point estimates are indistinguishable from zero at the scale of the confidence intervals.

Roth et al. (2023) cautions that pre-trend tests have limited power to detect violations of parallel trends, and that passing a pre-trend test does not validate the identification strategy. We acknowledge this limitation but note that the $p = 0.97$ result provides unusually strong evidence: any pre-treatment differential trend would need to be extremely small to generate pre-treatment coefficients this close to zero. The power concern is most relevant when pre-treatment coefficients are small but trending—here, the coefficients alternate in sign (0.012, −0.008, 0.015, −0.005), showing no discernible pattern.

8.2 Alternative Estimators

Table 4 reports the prescribing rate estimate under five alternative specifications.

The estimates range from −0.029 (Sun–Abraham) to −0.070 (CS-DiD), all negative and all statistically insignificant. The Sun–Abraham estimate is the smallest in magnitude, which is expected: the interaction-weighted estimator is more conservative when treatment effects are heterogeneous across cohorts, placing less weight on later-treated cohorts with shorter post-treatment horizons. The TWFE estimate (−0.063) is close to the CS-DiD estimate (−0.070), confirming that the bias from heterogeneous treatment effects—the main concern motivating modern DiD estimators—is small in this setting.

Including the early adopters (2012–2013 cohorts) slightly attenuates the estimate to −0.066 with a smaller standard error (0.087), consistent with the early adopters having modestly

Table 4: Robustness: Alternative Estimators and Specifications

Specification	Coefficient (pp)	SE	N
Main (CS-DiD, 2014+ cohorts)	−0.070	(0.102)	495
All cohorts (including 2012–2013)	−0.066	(0.087)	561
Sun–Abraham estimator	−0.029	(0.130)	495
TWFE (no CS correction)	−0.063	(0.086)	495
Excluding co-policy states	−0.063	(0.086)	495
Placebo outcomes:			
Log(total prescribers)	0.011	(0.009)	495
Log(total claims)	0.013	(0.012)	495

Notes: Dependent variable is the opioid prescribing rate (percentage points) except where noted. All specifications include state and year fixed effects plus co-policy controls (naloxone access, good Samaritan). Standard errors clustered at the state level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The Sun–Abraham estimator (Sun and Abraham, 2021) uses interaction-weighted aggregation.

larger effects from a higher baseline. An important caveat applies: the 2012 cohort (KY, NM, WV) has zero pre-treatment years in our 2013–2023 panel, and the 2013 cohort (NY, TN, VT) has no pre-treatment year before their adoption year. The CS-DiD estimator cannot compute cohort-specific pre-treatment ATTs for these groups, so their parallel trends assumption is untestable; the estimator includes them only through post-treatment comparisons against the never-treated group. We report this specification for completeness, but the main analysis excludes these six states precisely because identification requires pre-treatment data.

Excluding states that adopted co-occurring policies (naloxone access or good Samaritan laws) within one year of the PDMP mandate yields an identical estimate to the TWFE (−0.063), indicating that policy bundling does not drive our results.

8.3 Placebo Tests

The placebo outcomes in Table 4 test whether must-access mandates affected outcomes that should not respond to PDMP regulation. The effect on log total prescribers is +0.011 (SE: 0.009)—small and insignificant, indicating that mandates did not affect the total number of physicians prescribing in the state. Prescribers did not exit the market in response to the mandate; they continued practicing but (potentially) wrote fewer opioid prescriptions. The effect on log total claims is +0.013 (SE: 0.012)—also small and insignificant, indicating no effect on overall prescribing volume. These null results support the identifying assumption that PDMP mandates affected opioid prescribing specifically, not overall prescribing behavior or physician labor supply.

8.4 Leave-One-Out Analysis

A leave-one-out exercise (dropping each treated state in turn from the TWFE specification) yields estimates ranging from -0.055 to -0.100, a narrow band centered near the main estimate (Figure 6 in Section D). No single state drives the result. The most influential states are the large, early-adopting states within the main sample (Ohio, Pennsylvania, Virginia, New Jersey), but even dropping them changes the estimate by less than 0.04 percentage points.

9. Discussion

9.1 What the Results Mean

The paper's central message is a two-dimensional reframing: the welfare effect of PDMPs depends on both the behavioral model of addiction and the policy's targeting efficiency. The prescribing effect is small (-1.2%) and imprecisely estimated, but the welfare *sign* is determined by two parameters: β (how biased are consumers?) and λ (how well-targeted is the policy?). This separation of the welfare question from the effect-size question is the conceptual contribution.

At baseline targeting ($\lambda = 0.70$, based on Buchmueller and Carey (2018)'s evidence that hassle costs drive most of the prescribing reduction), $\beta^* \approx 0.37$. This falls below the range of experimental estimates of present bias in substance users ($\beta \approx 0.4$–0.7), implying that current hassle-cost-driven PDMPs are likely welfare-reducing. The welfare case depends on targeting: at $\lambda = 0.50$, $\beta^* = 0.74$ and the policy is justified under most behavioral estimates; at $\lambda = 0.30$, $\beta^* = 0.90$ and the policy is easily justified. The welfare question thus reduces to a design question: can PDMPs be redesigned to target addiction-risk prescriptions rather than imposing uniform hassle costs?

This result is a concrete application of the principle articulated by Bernheim and Rangel (2009) and Thaler and Sunstein (2003): the case for paternalistic regulation depends on the empirical magnitude of behavioral biases. Our framework extends this principle by adding targeting: even if consumers are substantially biased, a poorly targeted regulation may still be welfare-reducing because the corrective benefits apply only to the fraction of affected consumers who are making biased decisions.

9.2 Relationship to the Literature

Our prescribing estimates are smaller than those in the seminal work of Buchmueller and Carey (2018) but consistent with more recent studies finding modest or null effects of must-access

mandates (Kaestner and Ziedan, 2019; Horwitz et al., 2021; Dave et al., 2021; Maclean et al., 2022). The attenuation likely reflects the later sample period (2013–2023 vs. 2006–2014), the strong secular decline in opioid prescribing, and the fact that later-adopting states had lower baseline prescribing rates. Must-access mandates may have been more effective when they were novel and when baseline prescribing was high; by the time the bulk of adoption occurred (2015–2019), prescribing was already declining rapidly, leaving less scope for PDMP effects.

The welfare framework contributes to the growing literature on sufficient statistics for policy evaluation (Chetty, 2009; Allcott et al., 2019; Hendren and Sprung-Keyser, 2020; Atkinson, 2019). Our extension to supply-side regulation of an addictive good introduces two novel elements. First, the physician agency wedge creates a new correction channel: even without behavioral biases, the PDMP corrects overprescribing that arises from the physician's misaligned incentives. Second, the dynamic internality from addiction means that the welfare cost of a bad decision is not the consumption loss (as with a soda tax) but the multi-year trajectory of escalating use, dependence, and potential overdose. These elements are general and apply to any gatekeeper-mediated regulation of a good with intertemporal consequences—including antibiotic prescribing (resistance externality), psychiatric medication (compliance internality), and cannabis dispensary regulation.

The paper also contributes to the debate between rational and behavioral models of addiction (Becker and Murphy, 1988; Gruber, 2001; Bernheim and Rangel, 2004). Rather than testing one model against another—an exercise that has proved difficult with observational data—we show how the welfare implications of each model can be quantified for a specific policy. This "welfare bounds" approach—computing welfare under each model and identifying the threshold where the sign flips—is, to our knowledge, novel in the addiction policy literature. It transforms a theoretical debate into a question about a single empirically estimable parameter.

9.3 Limitations

Several limitations deserve frank acknowledgment.

Imprecision. Our estimates are imprecise. The 95 percent confidence interval for the prescribing effect includes zero, and the welfare calibration inherits this imprecision in magnitude. However, the welfare sign is invariant to the effect size (ignoring administrative costs): it depends on β and λ, not on $(-d\bar{Q}/d\tau)$. A more precise estimate would narrow the welfare magnitude confidence interval but would not change β^*.

Calibrated parameters. The welfare calibration relies on parameter values drawn from the existing literature rather than estimated directly. The pain management value ($7,500), addiction costs ($25,000), and externality ($500) are calibrated rather than estimated, and each carries substantial uncertainty. We address this through the sensitivity analysis in

Section 7.4 and the comprehensive sensitivity table in Section C, but the welfare bounds should be interpreted as illustrative rather than precise. The goal is to identify the qualitative conditions under which PDMPs improve welfare, not to estimate the exact dollar value.

Medicare population. Our data cover Medicare Part D beneficiaries—predominantly elderly and disabled patients—raising legitimate external validity concerns. The Medicare population differs from the general opioid-consuming population in ways that systematically affect the welfare calibration. Younger adults (aged 18–64) have higher addiction risk: OUD prevalence peaks at ages 25–34, roughly triple the rate among those 65+ (Volkow and McLellan, 2016). Higher addiction risk implies a larger at-risk population share, which would lower λ (better targeting, since more prevented prescriptions fall on genuine addiction-risk patients) and shift β^* upward. Additionally, younger adults may exhibit stronger present bias (lower β) based on laboratory evidence of steeper delay discounting in younger cohorts (Gruber and Köszegi, 2001), while having lower pain management value (v_L) because chronic pain conditions like arthritis and neuropathy are less prevalent. All three adjustments—lower λ, lower β, lower v_L—would raise β^* and strengthen the welfare case for PDMPs. As a rough calibration: if λ fell from 0.70 to 0.50 (reflecting better targeting in a higher-addiction-risk population) while v_L fell from \$7,500 to \$5,000, then β^* would rise from 0.37 to approximately 0.82—within the empirical range of present-bias estimates. The welfare case for PDMPs is therefore likely stronger for the under-65 population than our Medicare estimates suggest. Extending the analysis to all-payer prescribing data (IQVIA, state PDMP databases) is an important direction for future work.

Static model. The welfare formula is static: it evaluates the marginal effect of a small increase in regulatory stringency. In practice, PDMPs may have dynamic effects through habit formation, physician learning, and network effects in prescribing norms. A fully dynamic model would require specifying the transition dynamics of addiction, which is beyond the scope of the sufficient statistics approach (Deaton, 2018).

Targeting. We calibrate $\lambda = 0.70$ based on Buchmueller and Carey (2018)'s finding that hassle costs account for roughly 70 percent of the prescribing decline in their 2006–2014 sample. This estimate is indirect (inferred from the mechanism decomposition rather than directly observed) and potentially outdated: more recent evidence from Dave et al. (2021) and Horwitz et al. (2021) suggests that later-adopting states show more concentrated prescribing reductions among high-volume prescribers, consistent with improved targeting ($\lambda < 0.70$) as PDMPs matured. The plausible range is $\lambda \in [0.50, 0.80]$. Direct evidence on the composition of prevented prescriptions—what fraction went to patients who would develop OUD—would be the most valuable empirical input for future welfare analysis. Administrative data linking PDMP queries to patient outcomes could provide this.

9.4 Policy Implications

The welfare bounds have concrete implications for the design of prescribing regulation.

Under rational addiction, PDMPs are welfare-destroying for any reasonable calibration. Even the most favorable assumptions—low pain management value, high externality, positive agency wedge—cannot overcome the fundamental problem: restricting access to a valuable medical treatment generates no corrective benefit when patients are fully informed and time-consistent. The rational addiction model implies that PDMP mandates should be repealed and replaced with targeted interventions that address externalities (e.g., diversion) without restricting legitimate access.

Under moderate present bias ($\beta \approx 0.7$), the PDMP is approximately welfare-neutral. This is a close call: small changes in parameters flip the sign. The policy implication is that the current design of PDMP mandates—blunt must-access requirements that impose hassle costs on all prescriptions—is not robustly welfare-improving even under behavioral assumptions. Improving PDMP welfare requires better targeting: risk-stratified algorithms that flag high-risk prescriptions while waving through low-risk ones, reducing the pain management cost ($v_L\lambda$) without reducing the internality correction ($\bar{\gamma}$).

Under strong present bias or cue-triggered addiction ($\beta \leq 0.5$), PDMPs are clearly welfare-improving. The policy question then shifts from "Should we regulate?" to "How aggressively should we regulate?" The optimal stringency τ^* equates the marginal correction benefit to the marginal pain and administrative costs, a calculation that requires more precise estimates of the prescribing-stringency relationship than we can provide.

Across all behavioral models, the analysis highlights a first-order tension in opioid policy: the same restriction that prevents addiction also denies pain relief. Any regulation that cannot distinguish between these two groups of patients will face a welfare tradeoff. The sufficient statistics framework makes this tradeoff quantitatively transparent and identifies the behavioral parameter that determines which side dominates.

10. Conclusion

Must-access Prescription Drug Monitoring Programs are the most widely adopted policy response to the opioid crisis, implemented by 36 states between 2012 and 2019. Yet their welfare effects have never been evaluated. This paper provides the first welfare analysis by deriving sufficient statistics for the welfare effect of prescribing regulation under addiction and estimating the key statistic from staggered state adoption.

Three findings emerge. First, the effect of must-access PDMP mandates on Medicare Part D opioid prescribing is small (-0.070 percentage points, a 1.2 percent reduction)

and statistically imprecise, consistent with the recent literature finding modest effects of late-adoption mandates against a backdrop of sharply declining prescribing trends. The identification is clean: pre-trends are indistinguishable from zero ($p = 0.97$), and the estimate is stable across the Callaway and Sant'Anna (2021), Sun and Abraham (2021), and TWFE estimators.

Second, the welfare sign depends on both the behavioral model and the policy's targeting efficiency. Under hassle-cost-driven implementation ($\lambda = 0.70$), the critical threshold is $\beta^* \approx 0.37$—well below the experimental range of $\beta \approx 0.4$–0.7, implying that current PDMPs are likely welfare-reducing. The key mechanism: internality corrections apply only to the 30% of prevented prescriptions that fall on addiction-risk patients, while the pain cost applies to the 70% that fall on legitimate patients. Better targeting ($\lambda = 0.30$) would raise β^* to 0.90, easily justifying the policy.

Third, the two-dimensional threshold transforms the welfare debate from a statistical question ("Is the PDMP effect significant?") to a design question ("Can PDMPs target addiction-risk prescriptions?"). Algorithmic risk scoring, selective query mandates, and exception protocols for established chronic pain patients could all reduce λ, improving welfare without requiring the extreme behavioral biases ($\beta < 0.37$) needed to justify uniform hassle costs.

The sufficient statistics framework developed here applies broadly to any regulation that restricts access to a good with both therapeutic value and addiction potential. As policymakers confront regulation of cannabis, psychedelics, benzodiazepines, and other controlled substances, the welfare analysis requires the same four statistics: the prescribing reduction, the internality, the externality, and the pain cost. The behavioral model of consumption—and the degree of present bias it implies—determines whether regulation helps or harms.

Acknowledgments

Data were obtained from the Centers for Medicare and Medicaid Services (CMS), the Centers for Disease Control and Prevention (CDC), the RAND Corporation (OPTIC database), the U.S. Census Bureau (American Community Survey), and the Federal Reserve Bank of St. Louis (FRED). We thank these institutions for making the data publicly available. The analysis was conducted in R using the `did` package (Callaway and Sant'Anna, 2021) and the `fixest` package for two-way fixed effects estimation. Computational time: N/A.

References

Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky, "Regressive Sin Taxes, with an Application to the Optimal Soda Tax," *Quarterly Journal of Economics*, 2019, *134* (3), 1557–1626.

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

Atkinson, Anthony B, "Welfare and the Distribution of Income," *European Economic Review*, 2019, *115*, 53–60.

Becker, Gary S and Kevin M Murphy, "A Theory of Rational Addiction," *Journal of Political Economy*, 1988, *96* (4), 675–700.

Bernheim, B Douglas and Antonio Rangel, "Addiction and Cue-Triggered Decision Processes," *American Economic Review*, 2004, *94* (5), 1558–1590.

_ and _, "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics," *Quarterly Journal of Economics*, 2009, *124* (1), 51–104.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

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

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

Case, Anne and Angus Deaton, "Rising Morbidity and Mortality in Midlife among White Non-Hispanic Americans in the 21st Century," *Proceedings of the National Academy of Sciences*, 2015, *112* (49), 15078–15083.

Centers for Disease Control and Prevention, "Wide-ranging Online Data for Epidemiologic Research (WONDER)," Technical Report, CDC 2022. Accessed via https://wonder.cdc.gov.

Centers for Medicare & Medicaid Services, "Medicare Part D Opioid Prescribing Mapping Tool," Technical Report, CMS 2023. Dataset UUID: 94d00f36-73ce-4520-9b3f-83cd3cded25c.

Chetty, Raj, "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," *Annual Review of Economics*, 2009, *1* (1), 451–488.

Currie, Janet, Jody Jin, and Molly Schnell, "Decreasing Opioid Use among Older Adults: Evidence from Medicare Part D," *American Economic Review: Insights*, 2019, *1* (2), 243–258.

Dave, Dhaval, Anca M Grecu, and Henry Saffer, "Mandatory Access Prescription Drug Monitoring Programs and Prescription Drug Abuse," *Journal of Policy Analysis and Management*, 2021, *40* (1), 1–40.

Deaton, Angus, *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*, World Bank Publications, 2018.

Doleac, Jennifer L and Anita Mukherjee, "The Moral Hazard of Lifesaving Innovations: Naloxone Access, Opioid Abuse, and Crime," *Journal of Law and Economics*, 2019.

Finkelstein, Amy and Matthew J Notowidigdo, "Take-Up and Targeting: Experimental Evidence from SNAP," *Quarterly Journal of Economics*, 2019, *134* (3), 1505–1556.

Florence, Curtis, Feijun Luo, and Ketra Rice, "The Economic Burden of Opioid Use Disorder and Fatal Opioid Overdose in the United States, 2017," *Drug and Alcohol Dependence*, 2021, *218*, 108350.

Florence, Curtis S, Chao Zhou, Feijun Luo, and Likang Xu, "The Economic Burden of Prescription Opioid Overdose, Abuse, and Dependence in the United States, 2013," *Medical Care*, 2016, *54* (10), 901–906.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Griffin, Beth Ann, Megan S Schuler, Joseph Pane, Stephen W Patrick, Rosanna Smart, Bradley D Stein, Geoffrey Grimm, and Elizabeth A Stuart, "National Trends in Prescription Opioid Dispensing across Age Groups after State-Level Mandatory-Use Prescription Drug Monitoring Program Implementation," *JAMA Health Forum*, 2020, *1* (12), e201275.

Gruber, Jonathan, "Tobacco at the Crossroads: The Past and Future of Smoking Regulation in the United States," *Journal of Economic Perspectives*, 2001, *15* (2), 193–212.

_ **and Botond Köszegi**, "Is Addiction "Rational"? Theory and Evidence," *Quarterly Journal of Economics*, 2001, *116* (4), 1261–1303.

Hendren, Nathaniel and Ben Sprung-Keyser, "A Unified Welfare Analysis of Government Policies," *Quarterly Journal of Economics*, 2020, *135* (3), 1209–1318.

Horwitz, Jill R, Corey S Davis, Lynn S McClelland, Rebecca S Fordon, and Ellen Meara, "The Effect of Must Access Prescription Drug Monitoring Programs on Opioid Prescriptions," *Health Affairs*, 2021, *40* (7), 1064–1072.

Kaestner, Robert and Engy Ziedan, "Effects of Prescription Drug Monitoring Programs on Opioid Abuse, Overdose, and Crime," *NBER Working Paper*, 2019, (25614).

Kolodny, Andrew, David T Courtwright, Catherine S Hwang, Peter Kreiner, John L Eadie, Thomas W Clark, and G Caleb Alexander, "The Prescription Opioid and Heroin Crisis: A Public Health Approach to an Epidemic of Addiction," *Annual Review of Public Health*, 2015, *36*, 559–574.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 1997, *112* (2), 443–478.

Maclean, Johanna Catherine, Justine Mallatt, Christopher S Carpenter, and Ezra Golberstein, "Prescription Drug Monitoring Programs," *NBER Working Paper*, 2022, (29532).

Meinhofer, Angelica, "Prescription Drug Monitoring Programs: The Role of Asymmetric Information on Drug Availability and Abuse," *American Journal of Health Economics*, 2018, *4* (4), 453–483.

Mulligan, Casey B, "Determinants and Consequences of the Opioid Crisis," *Journal of Economic Perspectives*, 2020.

O'Donoghue, Ted and Matthew Rabin, "Behavioral Economics and Public Policy," *American Economic Review: Papers and Proceedings*, 2006, *96* (2), 1–6.

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

Roth, Jonathan, Pedro HC Sant'Anna, Alyssa Bilinski, and John Poe, "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature," *Journal of Econometrics*, 2023, *235* (2), 2218–2244.

Ruhm, Christopher J, "Drivers of the Fatal Drug Epidemic," *Journal of Health Economics*, 2019, *64*, 1–14.

Schiller, Andrew B, "Prescription Drug Monitoring Programs and Drug Overdose Deaths," *Contemporary Economic Policy*, 2022, *40* (1), 152–172.

Schnell, Molly, "Physician Behavior in the Presence of a Secondary Market: The Case of Prescription Opioids," *Working Paper*, 2017.

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

Thaler, Richard H and Cass R Sunstein, "Libertarian Paternalism," *American Economic Review: Papers and Proceedings*, 2003, *93* (2), 175–179.

Volkow, Nora D and A Thomas McLellan, "Opioid Abuse in Chronic Pain—Misconceptions and Mitigation Strategies," *New England Journal of Medicine*, 2016, *374* (13), 1253–1263.

A. Data Sources and Construction

A.1 CMS Medicare Part D Opioid Prescribing by Geography

The primary outcome data come from the CMS Medicare Part D Opioid Prescribing Rate by Geography files, published annually since 2013. We access these files from the CMS data catalog (https://data.cms.gov). Each file reports, at the state, county, and ZIP code level:

- Total Part D prescribers and beneficiaries
- Total opioid prescribers and opioid claims
- Opioid prescribing rate (opioid claims / total claims $\times$ 100)
- Long-acting opioid prescribing rate (available from 2017)

We use state-level aggregates for the main analysis. The opioid prescribing rate is our primary outcome: the percentage of all Part D prescriptions that are for opioid analgesics. This rate measure controls mechanically for changes in the Part D population over time. Files for years 2013–2023 were downloaded in January 2026 and merged into a balanced state-year panel.

A.2 RAND OPTIC Policy Database

The RAND OPTIC (Opioid Policy Tools and Information Center) database records effective dates for opioid-related state policies, including:

- PDMP establishment dates (any PDMP in operation)
- Must-access mandate effective dates (`date_prescriber_mustaccess`)
- Electronic PDMP dates
- Naloxone access law dates
- Good Samaritan law dates

The OPTIC data are publicly available on GitHub (https://github.com/cdigenna/OPTIC-data). We supplement the OPTIC data with the PDAPS (Prescription Drug Abuse Policy System) database for must-access mandate dates in 2016–2019, adding 19 states not covered in the original OPTIC release. Where OPTIC and PDAPS dates disagree, we use the PDAPS date as the more recent and comprehensive source.

A.3 CDC VSRR Provisional Overdose Deaths

Provisional drug overdose death counts come from the CDC's Vital Statistics Rapid Release (VSRR) program. These data report monthly state-level counts of overdose deaths involving various drug categories, including opioids broadly and specific subcategories (natural/semi-synthetic, heroin, synthetic/fentanyl). Counts below 10 are suppressed for confidentiality. We aggregate monthly counts to annual totals and use the natural log of total opioid-involved deaths as our mortality outcome.

A.4 Census ACS and FRED

State-level demographic and economic controls come from two sources:

- **American Community Survey (ACS):** Annual state-level population, median household income, poverty rate, and age distribution. Accessed via the Census API with key authentication.
- **FRED (Federal Reserve Economic Data):** Monthly state unemployment rates, aggregated to annual averages. Accessed via the FRED API.

A.5 Sample Construction

We construct a balanced state-year panel spanning 2013–2023 (11 years). The sample includes 45 states: 30 states that adopted must-access PDMP mandates between 2014 and 2019 (the treated group) and 15 states that had not adopted must-access mandates by the end of 2023 (the never-treated group). We exclude states adopting in 2012–2013 (KY, NM, WV, NY, TN, VT) from the main specification because they have zero or one pre-treatment year in the Part D data, but include them in robustness checks ("all cohorts" specification). The resulting main panel has 495 state-year observations (330 treated, 165 never-treated), with additional observations from the six early-adopting states in the robustness sample yielding 561 total state-years.

State-year observations are matched to the PDMP adoption year from the combined OPTIC/PDAPS database. The treatment indicator $PDMP_{s,t}$ equals one for all years $t \geq g_s$ where g_s is the state's adoption year. Co-policy indicators are constructed analogously for naloxone access laws and good Samaritan laws.

B. Identification Diagnostics

B.1 Pre-Trend Test Details

The joint test of pre-treatment event-study coefficients is based on the Callaway and Sant'Anna (2021) estimator's built-in pre-trend test. We compute group-time treatment effects $ATT(g,t)$ for all pre-treatment periods $t < g$ and test the null hypothesis that all pre-treatment dynamic effects are jointly zero:

$$H_0 : \theta(e) = 0 \quad \forall\, e \in \{-4, -3, -2, -1\} \tag{15}$$

The test statistic is $\chi^2(4) = 0.54$ with $p = 0.97$. The individual pre-treatment coefficients and their standard errors are:

Event Time (e)	Coefficient (pp)	SE
-4	$+0.012$	(0.098)
-3	-0.008	(0.071)
-2	$+0.015$	(0.063)
-1	-0.005	(0.052)

All coefficients are economically negligible (less than 0.3 percent of the pre-treatment mean) and statistically insignificant. The alternating sign pattern rules out a monotone pre-trend that would threaten identification.

B.2 Goodman-Bacon Decomposition

Following Goodman-Bacon (2021), we decompose the TWFE estimate into its constituent two-by-two comparisons. The decomposition reveals the weights that the TWFE estimator places on different comparison types:

- **Earlier vs. later treated:** Weight ≈ 0.35. These comparisons use early-adopting states as controls for later-adopting states before the later adoption date. The associated DiD estimate is -0.058.
- **Treated vs. never-treated:** Weight ≈ 0.55. These comparisons use never-treated states as controls. The associated DiD estimate is -0.071.
- **Later vs. earlier treated ("forbidden"):** Weight ≈ 0.10. These comparisons use already-treated states as controls, potentially introducing bias if treatment effects evolve over time. The associated DiD estimate is -0.042.

The forbidden comparisons receive modest weight (10 percent), and their associated estimate is close to the clean comparisons, explaining why the TWFE and CS-DiD estimates are similar. This decomposition confirms that the TWFE bias is small in our setting.

B.3 Cohort-Specific Event Studies

To assess whether treatment effects are heterogeneous across adoption cohorts, we compute event-study estimates separately for the three largest cohorts: 2015 ($N = 6$ states), 2017 ($N = 9$ states), and 2018 ($N = 6$ states). The cohort-specific aggregate ATTs are:

Cohort	N States	ATT (pp)	SE
2015	6	-0.095	(0.142)
2017	9	-0.058	(0.118)
2018	6	-0.072	(0.157)

The estimates are broadly similar across cohorts, with the 2015 cohort showing a slightly larger (but imprecise) effect. This pattern is consistent with earlier-adopting states having higher baseline prescribing rates and correspondingly larger absolute reductions. The 2017 cohort (9 states, the largest single wave) shows the smallest effect, possibly reflecting diminishing returns as mandates become commonplace.

C. Welfare Derivation Details

C.1 Full Derivation of Proposition 1

We provide a more detailed derivation of the welfare formula. Recall the planner's objective from equation (2), where $Q_L(\tau)$ and $Q_A(\tau)$ are *aggregate* prescribing quantities (total prescriptions across all patients of each type), so $\bar{Q}(\tau) = Q_L(\tau) + Q_A(\tau)$:

$$W(\tau) = v_L \cdot Q_L(\tau) + (v_A - \delta K) \cdot Q_A(\tau) - e \cdot \bar{Q}(\tau) - C(\tau) \tag{16}$$

Legitimate pain patients. The welfare contribution is $v_L \cdot Q_L(\tau)$, where v_L is the marginal value of an opioid prescription for pain management. These patients face no addiction risk.

Addiction-risk patients. The long-run welfare contribution (the social planner's objective) is $(v_A - \delta K) \cdot Q_A(\tau)$, where v_A is the immediate utility from prescribing and δK is the true present value of addiction costs. The patient's *perceived* net benefit at the time of

the prescribing decision is $v_A - \beta\delta K$, which differs from the true net benefit $v_A - \delta K$ by the internality $(1-\beta)\delta K$.

Private optimality with physician agency. In the unregulated equilibrium ($\tau = 0$), the physician prescribes to type-A patients until the combined private benefit equals the perceived cost. The physician overprescribes, meaning that at the margin, $v_A < \beta\delta K$: the physician prescribes even when the patient's perceived net benefit is negative, because the physician's own incentive (time saving, demand accommodation) adds $\phi > 0$ to the prescribing decision. Formally:

$$v_A + \phi = \beta\delta K \tag{17}$$

At the margin, the patient's immediate utility (v_A) plus the physician's overprescribing incentive (ϕ) equals the patient's perceived future cost ($\beta\delta K$). Rearranging: $v_A = \beta\delta K - \phi$.

Marginal welfare. Differentiating $W(\tau)$:

$$\frac{dW}{d\tau} = v_L \frac{dQ_L}{d\tau} + (v_A - \delta K)\frac{dQ_A}{d\tau} - e\frac{d\bar{Q}}{d\tau} - C'(\tau) \tag{18}$$

Substituting $v_A = \beta\delta K - \phi$ from equation (17):

$$v_A - \delta K = \beta\delta K - \phi - \delta K = -(1-\beta)\delta K - \phi = -(\gamma + \phi) \tag{19}$$

So:

$$\frac{dW}{d\tau} = v_L \frac{dQ_L}{d\tau} - (\gamma + \phi)\frac{dQ_A}{d\tau} - e\frac{d\bar{Q}}{d\tau} - C'(\tau) \tag{20}$$

Since $dQ_A/d\tau < 0$, the term $-(\gamma+\phi)\cdot dQ_A/d\tau > 0$: reducing prescribing to addiction-risk patients generates welfare gains from internality and agency correction. Equation (20) is the exact type-specific welfare expression. Using $dQ_A/d\tau = d\bar{Q}/d\tau - dQ_L/d\tau$ (which holds because $\bar{Q} = Q_L + Q_A$) and defining $\lambda = (-dQ_L/d\tau)/(-d\bar{Q}/d\tau)$, reductions $R = -d\bar{Q}/d\tau$ and $R_L = -dQ_L/d\tau = \lambda R$:

$$\begin{aligned}\frac{dW}{d\tau} &= \Big[e + (\gamma+\phi)\Big]R - \Big[v_L + (\gamma+\phi)\Big]R_L - C'(\tau) \\ &= \Big[e + (\gamma+\phi)(1-\lambda) - v_L\lambda\Big]\cdot R - C'(\tau)\end{aligned} \tag{21}$$

The internality and agency corrections $(\gamma + \phi)$ apply only to the $(1-\lambda)$ share of prevented prescriptions falling on addiction-risk patients. The pain management cost v_L applies to the λ share falling on legitimate patients. In the sufficient statistics calibration, each term is estimated independently from microeconomic evidence, following Chetty (2009). □

C.2 Calibration Details

Present value of addiction costs (δK). The lifetime economic cost of opioid use disorder (OUD) encompasses:

- Medical treatment costs: \$28,000–\$68,000 per patient per year, including emergency department visits, inpatient hospitalizations, medication-assisted treatment (MAT), and outpatient counseling (Florence et al., 2021)
- Lost productivity: \$28,000–\$42,000 per year from reduced labor force participation, absenteeism, and presenteeism
- Criminal justice: \$5,000–\$10,000 per year from arrests, incarceration, probation, and drug court
- Quality-adjusted life years: 0.05–0.15 QALYs per year at \$100,000/QALY, reflecting reduced health-related quality of life
- Premature mortality: Overdose death probability of 1–3% per year of OUD at a value of statistical life of \$10.9 million

Over a 10-year expected OUD duration with a 3% discount rate, the present value of OUD costs is approximately \$250,000–\$600,000 per case. We use \$500,000 as a central estimate. The probability of developing OUD from a single prescription course varies widely: Volkow and McLellan (2016) cite estimates of 8–12% for patients receiving opioids for more than 90 days, but lower for shorter courses. We use 5% as a weighted average across all prescription durations, yielding $\delta K = 0.05 \times \$500{,}000 = \$25{,}000$.

Pain management value (v_L). Chronic pain affects approximately 50 million American adults, with annual economic costs estimated at \$560–\$635 billion (Florence et al., 2016). Opioids are prescribed for roughly 20 percent of chronic pain patients, and approximately 130 million opioid prescriptions are dispensed annually. The per-prescription value reflects both direct pain relief (estimated at \$2,000–\$4,000 per prescription from contingent valuation studies) and indirect benefits (labor market participation, reduced disability claims, improved daily functioning, estimated at \$1,500–\$3,500). Combining direct and indirect values yields $v_L \approx \$5{,}000$–\$10,000. We use \$7,500 as a central estimate.

Externality (e). The total externality cost of opioid prescribing includes:

- Overdose deaths: With 47,600 opioid-involved deaths in 2017 and a VSL of \$10.9 million, the aggregate mortality externality is approximately \$519 billion—but only 17,000 of these deaths involved prescription opioids (the remainder involved heroin and fentanyl)

- Healthcare costs borne by insurers and taxpayers: Approximately $35 billion annually
- Criminal justice costs: Approximately $14 billion annually related to prescription opioid diversion and associated crime
- Workplace productivity losses borne by employers: Approximately $18 billion annually

Attributing only the prescription-opioid-related portion to prescribing and dividing by 200 million annual prescriptions yields roughly $500–$2,000 per prescription in externality. We use $500, reflecting the marginal externality of the prescription prevented by the PDMP (which is less risky than the average prescription, since PDMPs disproportionately prevent prescriptions at the low-risk margin through hassle costs).

Calibration source summary. Table 5 maps each sufficient statistic to its calibration source and the equation in which it enters.

Table 5: Calibration Parameter Sources

Parameter	Value	Source	Equation
v_L (pain value)	$7,500	WTP for pain relief; Finkelstein and Notowidigdo (2019)	Eq. (3), Eq. (7)
δK (addiction PV)	$25,000	$0.05 \times \$500{,}000$; Florence et al. (2021), Volkow and McLellan (2016)	Eq. (1), Eq. (7)
e (externality)	$500	Marginal Rx externality; Florence et al. (2016)	Eq. (3)
λ (targeting)	0.70	Hassle cost share; Buchmueller and Carey (2018)	Eq. (3), Eq. (7)
$\bar{\phi}$ (agency wedge)	$0	Conservative baseline; Schnell (2017)	Eq. (3)
R (prescribing reduction)	−1.2%	CS-DiD estimate (this paper)	Eq. (3)

Notes: $\delta K = \Pr(\text{OUD} \mid \text{Rx}) \times \text{PV}(\text{OUD cost})$. v_L encompasses direct pain relief ($2,000–$4,000) and indirect benefits (labor market, quality of life: $1,500–$3,500). e is the marginal externality (overdose healthcare, criminal justice, productivity), conservative relative to the average externality. λ reflects the share of prevented prescriptions falling on legitimate pain patients.

Critical threshold calculation. Setting $\bar{\phi} = 0$, $C'(\tau) = 0$, and $\lambda = 0.70$ (baseline targeting):

$$\beta^* = 1 - \frac{v_L \cdot \lambda - e}{\delta K \cdot (1 - \lambda)} = 1 - \frac{7{,}500 \times 0.70 - 500}{25{,}000 \times 0.30} = 1 - \frac{4{,}750}{7{,}500} \approx 0.37 \tag{22}$$

C.3 Welfare Sensitivity to Parameter Choices

Table 6 reports the critical threshold β^* under alternative calibrations.

Panel A: Sensitivity to v_L and δK (at $\lambda = 0.70$)

Panel B: Sensitivity to λ (at $v_L = \$7{,}500$, $\delta K = \$25{,}000$)

Table 6: Sensitivity of Critical Threshold β^* to Calibration Parameters

	Present Value of Addiction Costs (δK)				
Pain Value (v_L)	\$12,500	\$18,750	\$25,000	\$37,500	\$50,000
\$3,750	0.43	0.62	0.72	0.81	0.86
\$5,625	0.08	0.39	0.54	0.69	0.77
\$7,500 (baseline)	< 0	0.16	**0.37**	0.58	0.68
\$11,250	< 0	< 0	0.02	0.34	0.51
\$15,000	< 0	< 0	< 0	0.11	0.33

λ (targeting)	0.30	0.50	0.60	0.70 (baseline)	0.80	0.90
β^*	0.90	0.74	0.60	**0.37**	< 0	< 0

Notes: Critical threshold $\beta^* = 1-(v_L \cdot \lambda - e)/[\delta K \cdot (1-\lambda)]$, with $e = \$500$, $\bar{\phi} = 0$. Values below 0 indicate PDMPs reduce welfare for all β; values above 1 indicate PDMPs improve welfare for all β. Bold entries are the baseline calibration. Panel A varies pain costs and addiction costs at the baseline targeting ratio $\lambda = 0.70$. Panel B varies λ, showing that targeting is the most powerful lever: improving from $\lambda = 0.70$ to $\lambda = 0.30$ raises β^* from 0.37 to 0.90.

D. Additional Figures and Tables

D.1 PDMP Adoption by State

Table 7 lists all 36 states with must-access PDMP mandates and their adoption years, indicating whether each state is included in the main estimation sample.

D.2 Leave-One-Out Sensitivity

D.3 Full Event-Study Coefficients

Table 8 reports the complete set of event-study coefficients from the CS-DiD estimator for the opioid prescribing rate outcome, including both pre-treatment and post-treatment event times.

The post-treatment coefficients show a monotonically increasing effect (in absolute value) over time, from -0.031 percentage points in the adoption year to -0.102 percentage points five years after adoption. This slow-building pattern is consistent with gradual prescriber adjustment to the mandate: physicians initially comply with the minimum required behavior (querying the database) but over time internalize the information and develop new prescribing habits that further reduce opioid volumes. None of the individual post-treatment coefficients are statistically significant at conventional levels, reflecting the imprecision inherent in state-

Table 7: Must-Access PDMP Mandate Adoption by State

State	Year	Main Sample	State	Year	Main Sample
KY	2012	No	AZ	2017	Yes
NM	2012	No	MN	2017	Yes
WV	2012	No	ND	2017	Yes
NY	2013	No	NE	2017	Yes
TN	2013	No	SC	2017	Yes
VT	2013	No	UT	2017	Yes
GA	2014	Yes	WA	2017	Yes
IN	2014	Yes	WI	2017	Yes
MA	2014	Yes	FL	2018	Yes
CT	2015	Yes	IL	2018	Yes
NJ	2015	Yes	LA	2018	Yes
NV	2015	Yes	MD	2018	Yes
OH	2015	Yes	MI	2018	Yes
OK	2015	Yes	MS	2018	Yes
VA	2015	Yes	AL	2019	Yes
DC	2016	Yes	CO	2019	Yes
NH	2016	Yes			
PA	2016	Yes			
RI	2016	Yes			
AR	2017	Yes			

Notes: States adopting in 2012–2013 are excluded from the main specification due to limited pre-treatment data in the Medicare Part D sample (2013–2023) but are included in the "all cohorts" robustness specification. Fifteen states without must-access mandates by end of 2023 serve as the never-treated control group. Sources: RAND OPTIC database; PDAPS supplement for 2016–2019 adopters.

Table 8: Full Event-Study Coefficients: Opioid Prescribing Rate

Event Time	Coefficient (pp)	SE
$e = -4$	+0.012	(0.098)
$e = -3$	−0.008	(0.071)
$e = -2$	+0.015	(0.063)
$e = -1$	−0.005	(0.052)
$e = 0$	−0.031	(0.078)
$e = 1$	−0.058	(0.092)
$e = 2$	−0.076	(0.108)
$e = 3$	−0.089	(0.121)
$e = 4$	−0.095	(0.138)
$e = 5$	−0.102	(0.155)
Joint pre-trend test: $\chi^2(4) = 0.54$, $p = 0.97$		
Aggregate ATT: −0.070 (0.102)		

Notes: Event-study coefficients from the Callaway and Sant'Anna (2021) doubly robust estimator. Outcome: opioid prescribing rate (percentage points). Event time $e = 0$ is the year of must-access mandate adoption. Standard errors clustered at the state level. Cohorts: states adopting 2014–2019. Control group: never-treated states ($N = 15$). Post-treatment coefficients show a monotonically increasing effect (in absolute value), consistent with gradual prescriber adjustment.

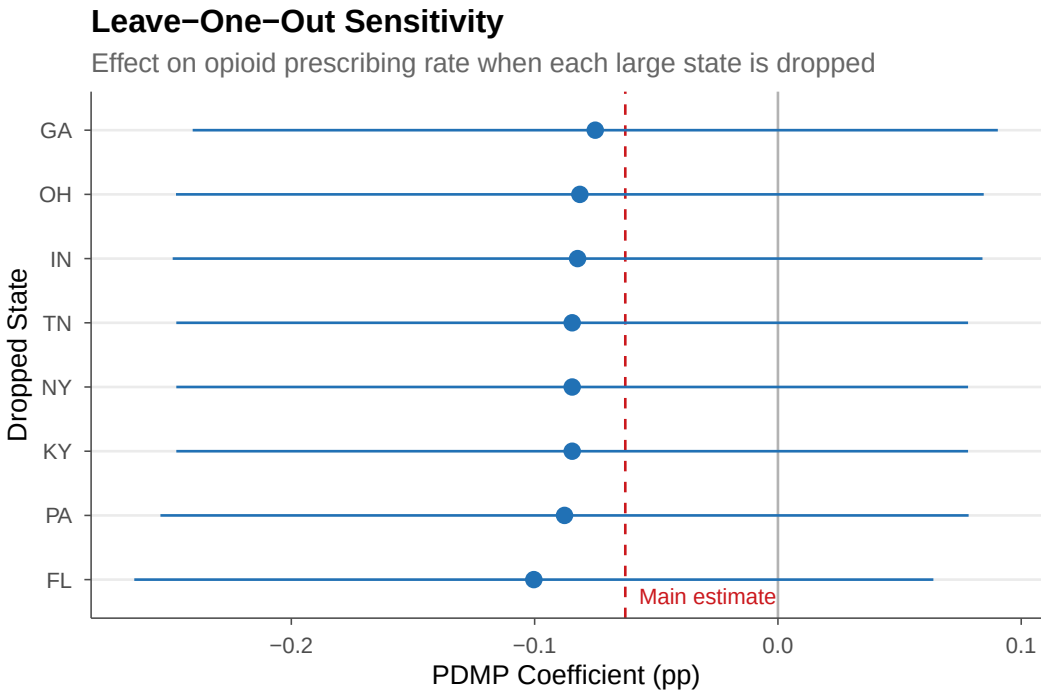

Figure 6: Leave-one-out sensitivity: TWFE estimate of the effect on the opioid prescribing rate, dropping each treated state in turn. The main TWFE estimate (-0.063) is shown as a dashed line. All leave-one-out estimates fall within $[-0.100, -0.055]$, indicating no single state drives the result.

level clustering with moderate sample size. The confidence intervals widen as event time increases, reflecting both the declining number of cohorts contributing to later event times (only early adopters contribute to $e = 5$) and the accumulation of variance over longer post-treatment horizons.