The Last Mile to First Treatment: Searching for Opioid Use Disorder Medication

By Lance Gui* 11/01/2025

[CLICK HERE FOR THE MOST RECENT VERSION]

Despite growing policy support for buprenorphine, the leading medication for opioid use disorder, treatment initiation remains low. Comprehensive insurance data from Washington State show that fewer than half of first-time patients fill their prescriptions, well below rates for other chronic conditions. Initiation is hindered not only by the physical cost of visiting pharmacies (e.g., distance) but also by uncertainty about availability. I estimate their joint effect using a structural sequential search model in which patients search across pharmacies with unknown availability. Normalizing the default (most frequently visited) pharmacy's search cost to zero, I identify search costs by comparing non-default choices (utility net of search costs) to default choices (utility only). I find that search costs account for 70% of observed treatment failures. Counterfactuals show that simple prescriber guidance on pharmacy availability could raise initiation share by 17%.

 $Keywords:\ treatment\ initiation,\ sequential\ search,\ pharmacy\ access$

JEL: L5, D83, I18

^{*} gui@arizona.edu. I am deeply grateful to Mo Xiao for her invaluable mentorship and guidance throughout this project. I thank Christian Cox, Hidehiko Ichimura, and Matthijs Wildenbeest for their thoughtful and detailed feedback. I also thank the Center for Management Innovations in Healthcare at the University of Arizona for funding the data purchase. I am especially grateful to Lorie Geryk, PhD, and Diwaker Tripathi, PhD, at the Washington State Health Care Authority for their assistance in accessing the Washington State All-Payer Claims Database. I appreciate the insightful discussions with Susan Bratton, M.D., Mark Grant, M.D., and Lynn Webster, M.D., whose expertise helped ensure the medical accuracy and real-world relevance of this paper. I also thank Aljoshcha Janssen, Ami Ichikawa, Austin Smith, Jingyuan Wang, Xuan Zhang, Yi Zhang, and seminar participants at the University of Arizona for their constructive feedback. All remaining errors are my own.

Conflict of Interests: I declare that I have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

I. Introduction

The opioid crisis remains the deadliest ongoing public health emergency in the United States. In 2024 alone, opioids caused 57,449 overdose deaths, bringing the 20-year cumulative toll to nearly one million. Among adults aged 16 to 65, opioids now account for 8.2% of all deaths. This exceeds deaths from suicide, firearm-related fatalities, and traffic accidents. Buprenorphine, a partial opioid agonist, is among the most effective treatments for opioid use disorder (OUD). By binding to the same receptors as opioids without producing the same euphoric high, it alleviates withdrawal symptoms and cravings, reducing the risk of relapse and overdose (Larochelle et al., 2018; Ma et al., 2019; Dever et al., 2024). Yet despite broad clinical and policy support, 87% of individuals with OUD remain completely untreated (Krawczyk et al., 2022).

This paper focuses on patients newly prescribed buprenorphine, for whom timely access is most critical. Delays at initiation often lead to a return to opioid use, derailing recovery before it begins. Yet after obtaining a prescription, over half of patients fail to fill it, a rate far below that for other chronic conditions. Why does treatment break down in the last mile?

Two reinforcing barriers contribute to the low rate of treatment initiation. First, buprenorphine remains difficult to access: only 57% of pharmacies nationwide stocked the medication in 2022 (Weiner et al., 2023). Second, patients face informational frictions: pharmacies are not required to disclose inventory and often decline to confirm buprenorphine availability by phone, citing concerns about diversion and theft. Even addiction specialists lack visibility into pharmacy inventories, limiting their ability to guide patients effectively. As one provider noted, "filling the prescription is entirely on the patient." These barriers were especially severe during the study period, when controlled substances were typically prescribed on paper rather than transmitted electronically. Patients had to visit pharmacies in person merely to learn whether the medication was available. Consequently, search costs include not only the physical effort of searching (e.g., travel) but also the uncertainty about pharmacy availability.

These barriers substantially raise the cost of initiating treatment, helping explain why so many patients fail to start. This naturally raises the policy question: how can initiation be improved?

¹Mortality counts are obtained from the CDC Wonder provisional multiple cause of death data [Accessed: 2025-07-18].

²In 2018, only 15.9% of controlled substance prescriptions in Washington State were transmitted electronically, according to Surescripts: Surescripts 2018 National Progress Report [Accessed: 2025-07-18].

One possibility is to reduce informational frictions by giving prescribers access to real-time pharmacy inventory data so they can direct patients to pharmacies that carry buprenorphine. Another is to offset high search costs with modest cash incentives, such as small gift cards. The effectiveness of such policies depends on disentangling patient preferences from the dual frictions of physical search costs and uncertainty about availability.

To empirically separate these factors, I combine detailed insurance data with a structural model of sequential search. The primary source is the Washington State All-Payer Claims Database (WA-APCD), covering 2014–2019, which links medical and pharmacy claims at the individual level across Medicaid, Medicare, and commercial insurers. Uniquely, this dataset includes patients with OUD, a population typically excluded from research due to federal privacy protections. From medical claims, I observe patients' diagnostic histories; from pharmacy claims, I observe drug use patterns.³ To capture pharmacy-level availability, I incorporate the Automated Reports and Consolidated Ordering System (ARCOS), which records shipments of buprenorphine to every retail pharmacy in the state. Linking ARCOS shipments to pharmacy claims allows me to infer pharmacy-level buprenorphine availability.

To analyze patients' search behavior, I build on Moraga-González, Sándor and Wildenbeest (2023), who estimate the dynamic sequential search model of Weitzman (1979) in a static discrete choice framework. I extend their approach with two modifications that account for pharmacy availability. First, I embed beliefs about pharmacy-level availability directly into physical search costs, constructing an effective search cost that scales inversely with perceived availability. For example, if a patient believes there is a 50% chance that a pharmacy has buprenorphine available, the effective cost of visiting that pharmacy doubles. Second, pharmacies that do not carry buprenorphine may be searched but cannot be purchased from; hence the model collapses to a discrete choice over pharmacies that carry buprenorphine.

A central challenge is distinguishing whether low initiation reflects weak patient preferences for buprenorphine or high effective search costs. To separately identify the two, I leverage a key empirical pattern: patients obtain nearly all prescriptions, controlled and non-controlled, from

³The data also include demographics such as age, sex, insurance type (Medicaid versus other), and residential ZIP code, which I use to measure travel distances to nearby pharmacies. To supplement individual characteristics, I merge in ZIP Code Tabulation Area–level data from the American Community Survey.

one default pharmacy. In the data, 96% of patients fill the majority of their prescriptions at a single location, and initiation share drop from 65% to 15% when this default does not carry buprenorphine. This sharp decline suggests that patients typically begin their search at the default and, when it is unavailable, often do not continue searching. Guided by these stylized facts, I assume in the model that patients always visit their default pharmacy. This is equivalent to assuming zero search cost at the default, so its utility reflects preferences alone, while non-default options embed both preferences and search costs. This structure enables identification using purchase data alone. Under these assumptions, the model reduces to a standard logit, with utility determined by preference and search cost parameters, allowing the same covariates to affect both.

The estimated model delivers two main findings. First, search is costly and therefore limited. Traveling one additional mile raises the effective search cost by \$0.62, about 3% of the average insured out-of-pocket cost for buprenorphine (\$21.33). While this may seem modest, it implies a much higher cost relative to price than in other markets.⁴ Consistent with high search costs, fewer than 40% of patients search more than once, relying instead on a single pharmacy even when it lacks the medication. In the model, these frictions explain 70% of treatment failures. Second, search costs and preferences exhibit sharp heterogeneity. Black patients, Medicaid and Medicare enrollees, and those with severe OUD are less likely to search beyond their default pharmacy and display lower estimated willingness to obtain buprenorphine, indicating disproportionate barriers for disadvantaged groups.

I use the estimated model to evaluate counterfactual policies that target different components of search frictions: physical search costs (e.g., traveling an extra mile) and perceived availability. The first set of policies reduce uncertainty alone. As a benchmark, I simulate a universal availability mandate requiring all pharmacies to carry buprenorphine. This eliminates uncertainty and expands the patient choice set. Under the status quo, I assume patients form beliefs based on the average availability rate among nearby pharmacies. When availability becomes known with certainty, these beliefs no longer inflate effective search costs, which are correspondingly

⁴For example, in the European car market, Moraga-González, Sándor and Wildenbeest (2023) estimate that traveling one kilometer imposes an average cost of €148. Given an average car price of €19,917, the search cost amounts to just 1% of the product's value per mile.

reduced. Treatment initiation improves from a baseline of 45% to 58% in this scenario. I then examine two more feasible information policies that avoid a mandate. Because pharmacies cannot publicly advertise buprenorphine availability, prescribers are natural intermediaries. Providing them with real-time inventory access allows patients to be directed to pharmacies with buprenorphine available, preventing wasted search and raising initiation by 17% relative to the status quo, an effect comparable to universal availability at much lower cost. Alternatively, prescribers may develop informal ties with a limited set of local pharmacies (e.g., five) that reliably carry buprenorphine. Knowing that such a subset has buprenorphine reduces uncertainty and lowers effective search costs, increasing initiation by 7%.

Next, I turn to mandatory electronic prescribing (e-prescriptions), now required nationwide, which reduces uncertainty but simultaneously raises physical search costs. By transmitting prescriptions electronically, e-prescribing removes the need for patients to visit pharmacies to check availability. Pharmacies can confirm inventory once the prescription is received, allowing patients to know whether buprenorphine is available before visiting. At the same time, however, e-prescriptions cannot be transferred without contacting the prescriber for a new one, raising the cost of switching if the pharmacy is unavailable. This rigidity raises search costs when patients attempt to switch pharmacies, discouraging continued search. Despite this additional burden, I find that the informational benefit of certainty dominates: treatment initiation improves under e-prescribing unless the effective search cost more than doubled, a threshold unlikely to be met.

Finally, I examine policies that offset search costs under electronic prescribing. Several states, including California and West Virginia, have piloted small incentives, typically \$10–\$20 gift cards, for Medicaid patients who reach treatment milestones (e.g., a negative urine test). Motivated by these initiatives, I consider whether incentives at the point of initiation could be effective. Because e-prescribing is now standard, I calibrate a scenario in which it reduces uncertainty but increases effective search costs by 20%, since prescriptions must be rewritten when pharmacies are unavailable. In this setting, a rebate of \$13.49 to Medicaid patients, who are the most price sensitive and least likely to initiate, raises the initiation share to 70%, a level comparable to that for other chronic conditions. Taken together, the findings highlight the value of a combined policy approach that addresses two components of search costs. With one policy

already in place and another piloted by several states, broader implementation appears feasible.

Related Literature: This paper proposes a novel approach to separately identify search costs and preferences in sequential search models for differentiated products without relying on directly observed search data.⁵ The key assumption is that one designated option, in this case the default, has zero search cost. The intuition is straightforward: when a consumer chooses a non-default option, it reveals that at least one additional search was undertaken. Although search is only partially inferred, such a choice indicates that the consumer was willing to incur a search cost. This approach proves valuable when direct search data are unavailable but understanding consumer behavior remains crucial. In school choice, students default to neighborhood public schools, yet many parents stick with these assignments despite better alternatives (Hastings, Kane and Staiger, 2006). Researchers typically cannot observe how many options parents consider or whether they search at all. By setting the default's search cost to zero, the model separately identifies preferences from search frictions. What appears as inertia can be attributed to parental inattention, but it may also reflect substantial search costs. This distinction matters for policy: eliminating the default assignment may prompt active choice, but it will not necessarily increase enrollment at higher-quality schools if search costs remain high.⁶

This paper contributes to a growing body of work evaluating policy responses to the opioid crisis. Most existing interventions have focused on restricting supply, such as the 2010 OxyContin reformulation (Severtson et al., 2013; Alpert, Powell and Pacula, 2018; Evans, Lieber and Power, 2018) and DEA crackdowns on rogue distributors (Soliman, 2023; Gui, 2024; Donahoe, 2025). While these efforts curtailed access to prescription opioids, they often failed to improve health outcomes and, in some cases, accelerated substitution into illicit markets. In contrast, demand-side policies have shown greater promise: expanding naloxone access reduces overdose mortality (Abouk, Pacula and Powell, 2019; Rees et al., 2017), and increasing the number of authorized buprenorphine prescribers improves treatment initiation (Gui, 2024). In this paper,

⁵For homogeneous products, search data is not required to identify search costs. For example, Hong and Shum (2006) attribute price dispersion solely to consumer search frictions and heterogeneity in search costs. Since all firms sell the same good, any price variation must be rationalized as an equilibrium outcome of search rather than quality or product differences.

⁶Similar dynamics arise in settings such as online shopping (e.g., defaulting to Amazon), voting (e.g., sticking to a party-line ballot), and retirement savings (e.g., auto-enrollment in investment plans). In all these cases, direct data on search behavior are often unavailable, yet search plays an important role in shaping choices.

I use availability at a patient's default pharmacy as an instrument for treatment initiation and find that initiation reduces ER visits by 20% relative to similar patients. This result offers causal confirmation of earlier observational findings (Sullivan, Szczesniak and Wojcik, 2021; Skains et al., 2023; Yarborough et al., 2024). More broadly, nearly all demand-side interventions depend on patients ultimately receiving medication for OUD. Reducing last-mile barriers can therefore amplify the effects of upstream policies aimed at expanding treatment initiation.

Second, this paper contributes to the empirical literature on estimating sequential search models, beginning with Weitzman (1979). It builds on recent advances by Moraga-González, Sándor and Wildenbeest (2023), who estimate the Weitzman model by drawing on insights from Armstrong (2017) and Choi, Dai and Kim (2018) to reformulate dynamic search as a static discrete choice problem. By imposing parametric restrictions on search cost distributions, they derive a tractable logit discrete choice model. This paper advances the literature by circumventing the need for direct search data. A key identifying assumption is that consumers always search their default option, allowing the model to separately recover preferences and search costs using observed purchase decisions alone. When such data are available (e.g., Kim, Albuquerque and Bronnenberg, 2010; Santos, Hortaçsu and Wildenbeest, 2012; Honka, 2014; Kim, Albuquerque and Bronnenberg, 2017; Ursu, 2018; Compiani et al., 2024), those moments can and should be incorporated. But in many markets, especially healthcare, no such data exist. In these cases, the proposed framework offers a tractable alternative. In addition, this paper extends the framework to account for product unavailability, a common feature in settings with supply disruptions, such as post-disaster access constraints (e.g., Luco, Klopack and Lewis, 2024). This extension preserves tractability: unavailable options are excluded from the choice set, and search costs are scaled by the perceived probability of availability.

Third, the paper contributes to the literature on default options in consumer choice, particularly in healthcare. Defaults carry significant weight in high-friction environments where consumers often deviate from the "best" option. As shown by Hortaçsu, Madanizadeh and Puller (2017) and Abaluck and Adams-Prassl (2021), reducing switching costs can improve welfare for consumers who remain anchored to suboptimal defaults. Models of inattention, such as those in Handel (2013), Ho, Hogan and Scott Morton (2017), Hortaçsu, Madanizadeh and Puller

(2017), Abaluck and Adams-Prassl (2021), and Heiss et al. (2021), typically sort consumers into two types: those who are inattentive and choose only the default or opt out, and those who are attentive and select from the full choice set. I propose an alternative model that embeds the default within a sequential search process. Departing from the default, meaning filling a prescription at a non-default pharmacy, requires incurring search costs. These costs are only worthwhile when the expected utility gain exceeds the cost of searching. This structure allows for intermediate behavior. Consumers may not be fully inattentive, but they also do not consider all available options. Instead, they evaluate a limited set of options shaped by the search process. The search-based view of default adherence offers a parsimonious, behaviorally grounded framework for modeling consumers who lie between fully attentive and fully inattentive, similar to the hybrid model proposed by Abaluck and Adams-Prassl (2021).

This paper proceeds as follows. Section II provides institutional background on OUD treatment and the limited availability of buprenorphine at pharmacies. Sections III and IV describe the data and highlight key empirical patterns. Section V presents a sequential search model under uncertain availability and outlines an estimation strategy that does not require search path data. Section VI reports the estimation results and discusses their implications. Section VII presents counterfactual policy analyses aimed at improving buprenorphine initiation. Section VIII concludes with a summary of findings and implications for policy. Additional empirical facts and details on data construction appear in the Appendix.

II. Institutional Background

This section provides institutional background on OUD treatment, with a focus on understanding the drivers of low buprenorphine prescription fulfillment. I highlight how pharmacy stocking decisions and limited information about availability contribute to treatment drop-off.

A. Medication for Opioid Use Disorder

"If you were to just imagine a medicine, or a chemical compound, that could stop the opioid epidemic, that medicine would probably look a lot like buprenorphine."

—Ethan Brook, The Atlantic

Buprenorphine is one of three medications approved by the U.S. Food and Drug Administration for treating OUD, alongside methodone and naltrexone. Among these, buprenorphine is distinguished by its safety profile and flexibility in administration, making it the most commonly prescribed treatment in office-based settings.

Methadone, a full opioid agonist, is subject to tight federal restrictions. It can only be dispensed through licensed opioid treatment programs, and patients are generally required to attend daily visits during the first 90 days of care. Naltrexone, an opioid antagonist, is more often prescribed for alcohol use disorder and requires complete detoxification prior to initiation, a barrier many patients are unable to overcome. In contrast, buprenorphine is a partial agonist with a ceiling effect that mitigates overdose risk while effectively suppressing cravings and withdrawal symptoms. Certified prescribers can offer it in office-based settings, expanding access beyond the confines of OTPs. During the study period, over 80% of pharmacologically treated OUD patients received buprenorphine.

The treatment pathway for buprenorphine resembles those of other long-term therapies for chronic conditions. Patients begin with an induction phase focused on clinical stabilization, followed by maintenance therapy that may extend for months or years. Behavioral therapy and peer support are often encouraged as complementary elements. Addiction specialists frequently liken buprenorphine to "insulin for OUD." It is not a cure, but a cornerstone of sustained disease management that enables individuals to return to functional daily life.

Initiating buprenorphine treatment is therefore a pivotal step in the recovery process. As with other chronic illnesses, early access and continuity of care are critical to achieving favorable outcomes. Yet, as this paper shows, frictions in access, especially at the pharmacy level, can obstruct treatment, even for patients who have already secured a valid prescription.

⁷Following regulatory changes during the COVID-19 pandemic, take-home methadone doses became more accessible. This flexibility has since been extended at the discretion of the prescribing practitioner. See SAMHSA guidance: Methadone Take-Home Flexibilities Extension.

⁸Buprenorphine is available in two formulations: one combined with naloxone (commonly branded as Suboxone) and one without. The combination is recommended as first-line treatment due to its deterrent effects on misuse. The mono-product is typically reserved for patients contraindicated for naloxone, such as pregnant women. Nearly all prescriptions in this study pertain to the combination formulation.

B. Pharmacy Stocking Behaviors

Despite buprenorphine's central role in treating OUD, its availability at retail pharmacies remains limited. Nationally, only 57.9% of pharmacies stocked the medication as of 2022 (Weiner et al., 2023). In Washington (2014–2019), only 56% of pharmacies that received any opioid shipments ever received buprenorphine. Given that its therapeutic role is often compared to insulin for diabetes, the reluctance of pharmacies to carry, let alone consistently stock, buprenorphine is striking.

The reasons pharmacies choose not to carry buprenorphine are not fully understood. A policy roundtable convened by the Substance Abuse and Mental Health Services Administration identified several commonly cited barriers. First, pharmacies may fear stigmatization, including concerns that stocking buprenorphine could attract clientele perceived as undesirable. Second, regulatory uncertainty, particularly regarding ordering thresholds, may discourage pharmacists from requesting inventory, out of fear of triggering compliance violations. Third, the economics of dispensing buprenorphine are often unfavorable. Pharmacists report high labor costs associated with verifying prescriptions and navigating insurance coverage, paired with low reimbursement rates. One widely circulated estimate suggests that pharmacies lose approximately \$10 per buprenorphine fill. 11

The challenge extends beyond whether a pharmacy ever stocks buprenorphine. Among those that do, inventory is often inconsistent or chronically insufficient. Because buprenorphine must be initiated within a narrow window after the onset of withdrawal, even short delays can disrupt treatment. As a result, patients are frequently forced to search across multiple pharmacies, introducing frictions that delay or deter initiation altogether.

⁹See: SAMHSA Policy Priority Roundtable on Buprenorphine Access in Pharmacies.

¹⁰The Drug Enforcement Administration shut down Oak Hill Hometown Pharmacy, citing an "imminent danger to public health and safety" due to its high volume of buprenorphine dispensing. Although a federal judge later ruled in the pharmacy's favor, noting it had acted professionally and within legal bounds, the DEA's action effectively ended its operations. Wholesale suppliers, wary of regulatory scrutiny, refused to resume shipments. See NPR Coverage [Accessed: 2025/07/19]

¹¹A natural question is why pharmacies continue to carry buprenorphine despite the apparent financial loss. Although the buprenorphine fill itself may be unprofitable, pharmacies can recoup losses through sales of other medications, particularly generics, as patients often have multiple prescriptions. In addition, pharmacies may fear losing the patient entirely if they do not carry buprenorphine. This behavior is evident in the data: about 40% of first-time patients switch all of their prescription fills to a pharmacy that carries buprenorphine, leaving their previous default pharmacy.

C. Frictions in Learning About Availability

Patients cannot easily resolve uncertainty about buprenorphine availability before visiting a pharmacy. Pharmacists are not legally required to disclose inventory levels, and many adopt a policy of nondisclosure for controlled substances. This practice, especially common in the early 2010s amid heightened fears of diversion, reflects concerns that confirming stock over the phone could attract individuals seeking opioids for non-medical use or increase the risk of theft.

These concerns appear in professional discourse. In a 2014 thread on *studentdoctor.net*, a forum used by pharmacists, contributors described common responses to telephone inquiries:

- "If it's a control, no matter what I say: 'out of stock.' I don't want any more of 'those' types of patients."
- "Just ask for the strength and quantity. Pretend to check. Then politely say no."

Many pharmacies will not even check inventory without first seeing a valid prescription. During the study period, most prescriptions for controlled substances were handwritten and physically carried by the patient. As of 2018, only 15.9% of controlled substance prescriptions were transmitted electronically. This meant that patients often had to visit pharmacies in person simply to determine whether the medication was available.

Even under electronic prescribing, where prescriptions can be transmitted before a patient arrives, new frictions emerge. If a pharmacy is out of stock, the prescriber must cancel the original and issue a new one; there is no mechanism for redirection. At the same time, federal regulations prohibit sending the same prescription to multiple pharmacies. In practice, this imposes added burdens on patients: when the default pharmacy lacks buprenorphine, each additional search requires contacting the prescriber to reroute the prescription again.

These institutional frictions have implications for the counterfactual scenarios considered later in the paper. Even if e-prescribing reduces uncertainty about inventory, it may raise the cost of switching pharmacies by requiring patients to contact their provider each time a redirection is needed. As a result, patients frequently initiate searches only after a failed fill attempt at their default pharmacy, and often without reliable information about which pharmacies are likely to carry buprenorphine.

III. Stylized Facts

This section presents empirical patterns that motivate the analysis. I first document the prevalence of default pharmacy behavior and show that buprenorphine initiation is low. Much of this shortfall is driven by lack of buprenorphine availability at patients' default pharmacies. Finally, I examine whether initiation shapes future use; starting encourages continuation, while non-initiation discourages it. Wherever appropriate, I benchmark buprenorphine against levothyroxine, the standard treatment for hypothyroidism. Levothyroxine is widely prescribed, routinely available, and not subject to the restrictions governing controlled substances. While both medications are first-line treatments for chronic conditions, a key distinction lies in their relative availability at pharmacies.

In Appendix B, I present three additional stylized facts. First, buprenorphine prescriptions are typically filled within a narrow therapeutic window, which highlights the importance of search for patients. Second, when the default pharmacy is unavailable, patients tend to substitute to pharmacies located either near their home ZIP code or near the default pharmacy, consistent with a sequential search process across pharmacies. Third, using default availability as an instrumental variable, I show that patients who initiate buprenorphine experience at least a 20% reduction in subsequent emergency room visits relative to those who do not, indicating that buprenorphine initiation is beneficial for patient health.

A. Default Pharmacy Behavior

I define a patient's default pharmacy as the pharmacy they visited most frequently in the year prior to their index diagnosis.¹³ Figure 1 shows that default pharmacy behavior is common for both buprenorphine and levothyroxine patients. Among those newly diagnosed with OUD or hypothyroidism, over 90% of patients filled at least half of their prescriptions at the same pharmacy prior to diagnosis. Specifically, 93% of buprenorphine patients and 88% of levothyroxine patients relied on a single pharmacy for the majority of their prior fills.

 $^{^{12}}$ Hypothyroidism, a condition in which the thyroid gland does not produce enough thyroid hormone, affects roughly 5% of the U.S. population.

¹³In rare cases, patients may split prescription volume evenly across multiple pharmacies. If so, I define the default pharmacy as the one nearest to the patient's home ZIP code. In addition, if a patient moves during the year, as indicated by a change in residential ZIP code, I assign the default based on the pharmacy most frequently visited prior to the move.




Figure 1. Concentration of Prior Fills at the Default Pharmacy

Note: This figure reports the share of past prescriptions filled at the default pharmacy for first-time buprenorphine and levothyroxine patients. Market share is defined as the fraction of all prescriptions filled at the patient's default pharmacy, which is the pharmacy they visited most frequently in the year prior to the relevant medical claim.

This consistent pattern across drug types suggests that default pharmacy behavior is not unique to OUD patients, but reflects a broader feature of prescription filling. While e-prescribing for controlled substances was limited during the study period, e-prescribing for non-controlled medications, such as levothyroxine, was already widespread. Most patients' prior prescriptions include at least some non-controlled substances, which are routinely transmitted electronically. These e-prescriptions are typically directed to the patient's preferred pharmacy on file, a default that remains fixed unless actively changed. E-prescribing platforms reinforce this behavior by auto-filling the last-used pharmacy, and medical guidelines encourage patients to consolidate prescriptions at a single location to reduce the risk of adverse drug interactions (Marcum et al., 2014). Therefore, patients tend to develop a default pharmacy. In what follows, I use the default pharmacy as the starting point for modeling patient search and decision-making under uncertainty.

B. Low Initiation in OUD Treatment


This section examines initiation of treatment among patients who have received a prescription.¹⁴ I classify patient behavior into three categories: filling the prescription at the default pharmacy, filling it at a non-default pharmacy, or not filling it at all.

Figure 2a displays the share of patients in each fulfillment category, conditional on whether buprenorphine was available at the patient's default pharmacy on the prescription date. I approximate daily inventory levels by linking pharmacy claims to ARCOS shipment data (see Appendix A.A2 for details). Three key patterns emerge. First, although most patients had default pharmacies with buprenorphine available, the majority of non-fills occurred when the drug was unavailable at the default. Second, availability at the default pharmacy significantly increased initiation: only 30% of patients failed to fill the prescription when buprenorphine was available. Third, when the default pharmacy lacked the medication, just 15% of patients filled the prescription elsewhere. This suggests that most do not search extensively or are unable to locate an alternative with supply.¹⁵


Figure 2b presents unconditional purchase patterns for buprenorphine and levothyroxine patients, that is, patterns not conditioned on drug availability. About 30% of hypothyroidism patients fail to fill a prescription following diagnosis, compared to nearly 55% of OUD patients. Among levothyroxine users, fills occur at the default pharmacy 13 times more often than at non-default locations. For buprenorphine, that ratio falls to 4:1, indicating greater deviation from the default, either due to increased search effort or lower availability at the usual pharmacy. Crucially, when restricting to buprenorphine patients whose default pharmacy has the drug available (in the prior figure), their fill rates and default usage closely resemble those of levothyroxine patients. This pattern suggests that low buprenorphine initiation is not driven primarily by patient non-compliance but rather by limited availability and the resulting search frictions.

¹⁴This inference is based on medical claims data, which do not specify the exact drug prescribed. However, both hypothyroidism and OUD are life-threatening without medication, making it highly likely that treatment was prescribed. Among patients who had a follow-up visit with the same provider, only 3% failed to generate any pharmacy claims, further supporting the presumption of prescribing.

¹⁵For this group of patients, I examine subsequent purchase patterns. As shown in Appendix B, patients tend to purchase from pharmacies located either near their home ZIP code or their default pharmacy. This empirical pattern motivates the structure of the search model.



(a) Purchase Location Conditional on Default Pharmacy Availability



(b) Purchase Location Following Initial Diagnosis

Figure 2. Pharmacy Purchase Patterns for New Patients

Note: This figure illustrates where first-time patients fill their prescriptions: either at their default pharmacy, at a non-default pharmacy, or not at all. Panel (a) conditions on the availability of buprenorphine at the default pharmacy, defined as having sufficient inventory to fill the prescription on the date it was written. Panel (b) shows the unconditional purchase shares for patients prescribed buprenorphine (used to treat OUD) and levothyroxine (a common and widely available treatment for hypothyroidism).

C. Initiation Shapes Future Use




Figure 3. Monthly Buprenorphine Fills Following Initial Diagnosis

Note: This figure plots the monthly share of patients filling at least one buprenorphine prescription in the 12 months after their first OUD-related doctor visit. Results are shown separately for patients who filled their initial prescription and those who did not.

I focus on first-time patients diagnosed with OUD for two reasons. First, search frictions and pharmacy availability are most salient at treatment initiation, before patients have established a pharmacy relationship. In contrast, refills can often be coordinated in advance, even when a pharmacy is temporarily out of stock, and patients can refill before their current medication runs out. Second, the first filled prescription is a decisive margin: patients who begin treatment are far more likely to remain engaged, while those who do not initiate rarely start later.

Figure 3 highlights the pivotal role of initiation. Among patients who filled their initial prescription, nearly 60% remained on buprenorphine in the following month, and more than 20% continued treatment one year later. By contrast, fewer than 3% of patients who failed to

fill their initial prescription ever initiated treatment within the following year. The gap is not merely large in levels; it is persistent across time. Every subsequent month shows markedly higher treatment rates among initiators compared with non-initiators. This pattern underscores a central fact: taking buprenorphine for the first time has enduring consequences. Initiation substantially raises the likelihood of continued treatment, while failing to initiate effectively forecloses future engagement.

IV. Data

To study pharmacy search behavior and buprenorphine access in Washington State, I use three primary datasets. This section summarizes the key features of each dataset relevant to the analysis. Table E1 reports summary statistics for the main variables used.

A. Washington State All-Payer Claims Database

The WA-APCD includes medical and pharmacy claims for most insured residents of Washington State, covering public payers (Medicaid and Medicare) as well as a subset of private plans. It provides detailed claim-level information on patient demographics, diagnoses, procedures, provider identifiers, payment amounts, and residential ZIP codes. As a patient-centered dataset, the WA-APCD allows for the reconstruction of individual medical histories, such as prior diagnoses of substance use disorder and prior-quarter opioid consumption, and enables linkage between medical visits and subsequent pharmacy fills.

A unique feature of this dataset is that it includes patients diagnosed with OUD, a subset of substance use disorders historically excluded from administrative claims data due to federal privacy protections. To my knowledge, this is the first paper to use detailed patient-level insurance data to study treatment-seeking behavior among individuals with an OUD diagnosis.

I restrict the sample to patients who received their first OUD diagnosis between 2014 and 2019. Using medical claims, I identify the initial OUD-related visit and then track whether the patient filled a buprenorphine prescription, and at which pharmacy, using pharmacy claims. Further details on sample construction are provided in Appendix A.A1.

B. Automated Reports and Consolidated Ordering System

ARCOS is a national surveillance system maintained by the Drug Enforcement Administration under the Controlled Substances Act of 1971. ARCOS tracks opioid shipments across the supply chain, from manufacturers and distributors to hospitals and pharmacies. For each transaction, the dataset records the shipment date, quantity, dosage, National Drug Code, and the identities of sender and recipient.

ARCOS plays a central role in this paper by enabling the construction of daily pharmacy-level inventory for buprenorphine. Pharmacy claims contain information on drug formulation, strength, and fill date, which I use to measure daily outflows. Shipment data from ARCOS provide the corresponding inflows. By aggregating shipments and claims over time and subtracting cumulative claims from cumulative shipments, I estimate whether a pharmacy had buprenorphine available on a given day, including the date of a patient's prescription. This inventory measure is not without limitations. In particular, not all buprenorphine is dispensed through insurance claims, and I abstract from formulation differences by focusing on the presence of buprenorphine as the active ingredient. Construction details and a discussion of caveats appear in Appendix A.A2.

C. Other Data Sources

ZCTA-Level Demographics. Because individual-level demographic data in the WA-APCD are limited, I supplement the analysis with ZIP Code Tabulation Area-level characteristics from the American Community Survey 5-year estimates. These include median income, educational attainment (above high school), poverty rate, and unemployment rate.

Provider Characteristics. I use the National Plan and Provider Enumeration System (NPPES) to obtain information on prescribers, including specialty, gender, and years of experience. Historical NPPES snapshots were accessed through the NBER data repository. ¹⁶

Product Classification. To distinguish buprenorphine formulations intended for OUD from those used for pain management, I rely on data scraped from the DailyMed website. By querying

¹⁶Link to NBER collection [Accessed: 2025-01-15].

National Drug Codes for buprenorphine products, I obtain detailed product descriptions and labeled indications for use.

V. Model

This section develops a sequential search model to describe how first-time buprenorphine patients search across pharmacies for treatment. The framework builds on Weitzman (1979)'s optimal sequential search strategy. I extend it to account for cases where a search attempt yields a reward equal to zero, specifically when the patient discovers that a pharmacy does not have buprenorphine available.¹⁷ Building on this foundation, the estimation approach follows Moraga-González, Sándor and Wildenbeest (2023), which applies insights from Armstrong (2017) and Choi, Dai and Kim (2018) to reformulate sequential search as a static discrete choice problem. In this framework, consumers select the option that maximizes the minimum of their reservation value and realized utility across all available alternatives. Importantly, the model does not require auxiliary data on search sequences or the number of pharmacies visited. Identification hinges on asymmetries in the utility specification between default and non-default options. Specifically, the utility of selecting the default pharmacy is not subject to search costs, whereas all non-default options incur such costs. This asymmetry generates variation in observed choices that can be exploited for estimation.

A. Utility Specification

I consider a market with J different pharmacies, indexed by j = 1, 2, ..., J. The indirect utility that consumer i derives from visiting pharmacy j and obtaining buprenorphine is given by:

(1)
$$u_{ij} = \beta_0 + \alpha_{ins} P_{ij} + X_{ij} \beta + \epsilon_{ij},$$

where P_{ij} is the price paid by consumer i at pharmacy j and for the price coefficient α_{ins} , The price sensitivity parameter, α_{ins} , is allowed to vary by insurance type (Medicaid, Medicare,

¹⁷This is discussed in the Application section of Weitzman (1979), where box i contains one of two outcomes: either zero reward ("failure") with probability $1 - p_i$, or a positive reward ("success") with probability p_i .

or commercial coverage).¹⁸ The vector X_{ij} includes consumer, provider, and pharmacy characteristics: demographics (e.g., age, gender, race), health indicators (e.g., past-quarter opioid consumption, prior diagnoses of substance use disorder or mental illness), physician characteristics (e.g., addiction specialization, years of experience), pharmacy characteristics (e.g., chain affiliation, same-chain indicator with the default pharmacy), and insurance type (e.g., Medicaid, Medicare, or commercial coverage).

The idiosyncratic term ϵ_{ij} is assumed to be *i.i.d.* across pharmacies and to follow a Type I Extreme Value distribution. It represents a match-specific parameter that captures the "fit" between consumer i and pharmacy j. While the drug itself is homogeneous, the match value ϵ_{ij} varies across pharmacies and patients. One interpretation is that it reflects factors only revealed upon visiting a pharmacy, such as the quality of pharmacist interactions or wait times. These interactional frictions are especially relevant in OUD treatment, where stigma remains a substantial barrier to care. For example, if a pharmacist treats a patient dismissively or with hostility, the patient may simply leave without filling the prescription. The utility of not filling a buprenorphine prescription is denoted by $u_{i0} = \epsilon_{i0}$, where j = 0 represents the outside option.

Equation (1) applies only to pharmacies where buprenorphine is available. For pharmacies without buprenorphine, I define the realized utility as $u_{ij} = u_{i0} - \iota$, where ι is an arbitrarily small positive value. This ensures that unavailable pharmacies are strictly dominated by the outside option and thus never chosen.

Consumers must visit a pharmacy to learn the realization of ϵ_{ij} and thus their true utility. I model any trip beyond the default pharmacy as a discrete search attempt, originating from either the patient's home ZIP code or the default pharmacy. Search is sequential and ordered, with costless recall. After visiting a pharmacy, the consumer decides whether to purchase, continue searching, or opt for the outside option. Before initiating search, the consumer knows the location of each pharmacy, the utility of the outside option, and the distribution F from which ϵ_{ij} is drawn.

¹⁸Prices in this setting are largely determined by insurance contracts and network agreements, limiting pharmacies' ability to adjust prices in response to consumer search. In contrast to standard search markets where firms may exploit inattention through markups (Moraga-González, Sándor and Wildenbeest, 2023), price endogeneity is not a primary concern here.

¹⁹I remain agnostic about whether the search originates from home or from the default pharmacy; both distances are included as covariates in the search cost function.

Let \tilde{c}_{ij} denote the effective search cost in conjunction with patient i's belief on availability for patient i when visiting pharmacy j. Effective search costs vary across consumers and pharmacies. I assume that the cumulative distribution of effective search costs, $F_{ij}^{\tilde{c}}$, depends on search cost covariates through a location parameter μ_{ij} . The effective search cost distribution has full support, though only its non-negative portion affects consumer behavior.

B. Optimal Sequential Search

I characterize optimal consumer search behavior using a modified version of the Weitzman rule. Let $\phi_{ij} \in (0, 1]$ denote patient *i*'s belief about the probability that pharmacy *j* has buprenorphine available.²⁰ In this framework, consumers decide whether to search a given option by comparing the expected benefit of searching to the cost of search. In my context, the expected benefit of searching pharmacy *j* depends on both the likelihood of availability and the potential gain in utility if buprenorphine is indeed available.

Let r denote the best utility the consumer has observed so far. The expected gain from searching pharmacy j is:

(2)
$$\phi_{ij} \underbrace{\int_{r}^{\infty} (z - r) dF_{ij}(z)}_{\equiv H_{ij}(r)},$$

where $F_{ij}(z)$ is the cumulative distribution function of the realized utility when pharmacy j has buprenorphine available.

The expected gain has two components: the probability of availability, ϕ_{ij} , and the expected utility gain conditional on availability, denoted $H_{ij}(r)$. By assumption, if a pharmacy does not have buprenorphine, its realized utility is set to be infinitesimally lower than the utility of the outside option. Thus, the expected gain from searching an unavailable pharmacy is effectively zero. This justifies modeling the total expected gain as $\phi_{ij}H_{ij}(r)$, with the zero-utility outcome implicitly folded into the belief term ϕ_{ij} .

 $^{^{20}}$ The fact that $\phi_{ij}\neq 0$ arises mechanically from the model's construction. Interpreting $\phi_{ij}=0$ as full certainty that pharmacy j does not carry buprenorphine implies that the patient would never search that pharmacy. Consequently, the estimation remains unchanged even when $\phi_{ij}=0$ for a subset of pharmacies.

If consumer i's expected gains from searching pharmacy j exceed the cost c_{ij} , the consumer will choose to search. I define the corresponding reservation value r_{ij} as the solution to the following equation:

$$\phi_{ij}H_{ij}(r_{ij}) - c_{ij} = 0.$$

In addition, I define the effective search cost as:

(4)
$$\tilde{c}_{ij} \equiv \frac{c_{ij}}{\phi_{ij}},$$

The term c_{ij} reflects the physical or cognitive cost of visiting pharmacy j, independent of availability beliefs. Equation (4) highlights a key insight: lower expected availability inflates the effective cost of search. For example, if $\phi_{ij} = 0.5$, the effective cost of visiting pharmacy j is doubled. This framework generalizes earlier models in the literature, which implicitly assume $\phi_{ij} = 1$ and estimate c_{ij} directly.

I estimate \tilde{c}_{ij} as a function of observable search cost covariates, including distance from home or from the default pharmacy, pharmacy characteristics (such as chain affiliation), and patient demographics. A key advantage of this approach is that it imposes no restriction on the correlation between physical search costs c_{ij} and availability beliefs ϕ_{ij} , including in the unobservables. The limitation is that the two components are not separately identified. For instance, a patient may avoid a chain pharmacy such as CVS either because verification is time-consuming (high c_{ij}) or because they believe the pharmacy is unlikely to carry buprenorphine (low ϕ_{ij}). While the individual components cannot be separately recovered, their combined effect on behavior is fully captured by \tilde{c}_{ij} .²¹

Importantly, since H_{ij} is decreasing and strictly convex, Equation (3) has a unique solution. Therefore, $r_{ij} = H_{ij}^{-1}(\tilde{c}_{ij})$. Following Lemma 1 from Moraga-González, Sándor and Wildenbeest

 $^{^{21}}$ One alternative is to impose a scalar structure on ϕ_{ij} and directly assume the parametric distribution for c_{ij} , which implicitly requires that the unobservables in c_{ij} and ϕ_{ij} are uncorrelated. This restriction is difficult to justify. Another possibility would be to use settings with variation in beliefs, such as comparing handwritten prescription periods with e-prescription periods. Under e-prescriptions, patients learn availability before visiting, which would aid separate identification of ϕ_{ij} . However, such data are not available in this setting.

(2023), the reservation value r_{ij} can be expressed as:

(5)
$$r_{ij} = \delta_{ij} + H_0^{-1}(\tilde{c}_{ij})$$

where $H_0(r) = \int_r^{\infty} (z - r) dF(z)$, and δ_{ij} is the deterministic component of utility from Equation (1). This expression makes clear that the reservation value depends on both the expected utility from choosing a given pharmacy and the costliness of searching it. Specifically, higher effective search cost \tilde{c}_{ij} lowers the reservation value, making consumers more selective in which pharmacies to visit. Conversely, higher observable match quality δ_{ij} lowers the bar for search.

C. Discrete Choice Problem

Armstrong (2017) and Choi, Dai and Kim (2018) show that the optimal sequential search strategy can be equivalently expressed as a max-min decision rule: the consumer selects the option that maximizes the minimum of the reservation value and the realized utility. This framework can be extended to incorporate product availability. In my setting, consumers visit pharmacies in descending order of reservation values. At each pharmacy, they observe whether buprenorphine is available and, if so, draw the realized utility. If the drug is unavailable, the realized utility is strictly lower than that of the outside option. Search continues until the realized utility at a visited pharmacy exceeds the reservation value of the next-best unvisited option. A formal proof of this equivalence is provided in Appendix C.

A key feature of the model is that pharmacies without buprenorphine are excluded from the patient's choice set because they cannot dispense the medication. This exclusion reflects both institutional constraints and modeling assumptions. Patients cannot purchase buprenorphine from a pharmacy that does not have it, and the realized utility from such pharmacies is assumed to be strictly lower than that of the outside option. Accordingly, the discrete choice problem is defined over the set of pharmacies with buprenorphine available to consumer i, denoted A_i . For each available pharmacy $j \in A_i$, I define the latent utility w_{ij} as the minimum of the reservation

value r_{ij} and the realized utility u_{ij} :

$$(6) w_{ij} = \min\{r_{ij}, u_{ij}\}.$$

To derive a tractable expression for the distribution of w_{ij} , I adopt a specific parametric form for the distribution of effective search costs. I assume:

(7)
$$F^{\tilde{c}}(\tilde{c}) = \frac{1 - \exp\left(-\exp\left(-H_0^{-1}(\tilde{c}) - \mu_{ij}\right)\right)}{1 - \exp\left(-\exp\left(-H_0^{-1}(\tilde{c})\right)\right)},$$

where μ_{ij} is the location parameter for the search cost distribution.

Under this specification, the induced distribution of w_{ij} follows a Type I Extreme Value distribution with location parameter $\delta_{ij} - \mu_{ij}$:

(8)
$$F_{ij}^{w}(w) = \exp(-\exp(-(w - (\delta_{ij} - \mu_{ij})))).$$

This closed-form result relies on the particular choice of the effective search cost distribution and is formally derived in Proposition 1 of Moraga-González, Sándor and Wildenbeest (2023).

Then, we can specify the latent utility of pharmacy j for patient i as:

(9)
$$w_{ij} = (1 - D_{ij})(\delta_{ij} - \mu_{ij}) + D_{ij}\delta_{ij} + \xi_{ij},$$

(10)
$$w_{i0} = \xi_{i0}$$
,

where D_{ij} is an indicator equal to one if pharmacy j is patient i's default pharmacy. δ_{ij} is the deterministic component of utility from Equation (1), and μ_{ij} is the location parameter for the search cost distribution. A positive value of μ_{ij} is required for $F^{\tilde{c}}(\tilde{c})$ to be a proper distribution. I parameterize μ_{ij} using a log-exp functional form:

(11)
$$\mu_{ij} = \log(1 + \exp(\lambda_0 + Z_{ij}\lambda)),$$

 $^{^{22}}$ Each patient has exactly one default pharmacy, such that $\sum_{j} D_{ij} = 1$. While the framework could be extended to allow multiple defaults or no default at all, I restrict attention to a single default pharmacy in order to define distance from the default in a consistent way.

where Z_{ij} includes observed search cost shifters that vary by consumer and/or pharmacy. These may include distance from the patient's home ZIP code, patient demographics, and pharmacy chain affiliations.

The conditional choice probability that patient i selects pharmacy j is:

(12)
$$s_{ij} = \frac{\exp(\kappa_{ij})}{1 + \sum_{j \in A_t} \exp(\kappa_{ij})},$$

where $\kappa_{ij} = (1 - D_{ij})(\delta_{ij} - \mu_{ij}) + D_{ij}\delta_{ij}$. As discussed earlier, pharmacies without buprenorphine are never chosen and are excluded from estimation.

In sum, this discrete choice framework yields a closed-form logit model, despite the underlying sequential search structure. Availability enters as a constraint on the choice set, and search costs are embedded in the utility index via μ_{ij} . The model is computationally tractable and does not require auxiliary data on search sequences or moments that separately identify preferences and search costs.

D. Estimation and Identification

I estimate the model via maximum likelihood. The log-likelihood function is:

(13)
$$\mathcal{L}(\theta) = \sum_{i=1}^{N} \sum_{j \in \{A_{it}, 0\}} Y_{ij} \log(s_{ij}),$$

where $Y_{ij} = 1$ if patient *i* fills their buprenorphine prescription at pharmacy *j*, and $Y_{ij} = 0$ otherwise. The parameter vector θ includes the price coefficient α , the preference parameters β , and the search cost parameters λ .

To reduce computational burden, I restrict each consumer's choice set to pharmacies within 25 miles for urban ZIP codes and 50 miles for rural ZIP codes. This restriction not only accelerates estimation but also helps remove problematic observations, such as patients who may have moved but failed to update their recorded ZIP code.

Identification: Identification boils down to separating the components of utility, δ_{ij} , from those of search cost, μ_{ij} . The model is identified through two main sources of variation.

First, I exploit asymmetries in how w_{ij} is constructed for default and non-default options. For the default pharmacy, I assume that patients always search it which implies that its effective search cost is zero. As a result, the default's utility depends only on δ_{ij} . In contrast, the utility of non-default pharmacies reflects both δ_{ij} and μ_{ij} . Variation in observed choices between the default and the outside option (as it is normalized to zero) identifies δ_{ij} . Given the combined effect of $\delta_{ij} - \mu_{ij}$ are identified through variation in observed choices like a standard logit, I can identify μ_{ij} through differencing.

For instance, consider the constant term: let β_0 be the coefficient on the constant in δ_{ij} , and λ_0 the corresponding coefficient in μ_{ij} . A change in β_0 affects the utility of default options but leaves the utility of the outside option unchanged. By comparing the default option to the outside option, we can identify β_0 . Given that the total effect of the constant is identified, we can then back out λ_0 by differencing. This logic generalizes to any covariate that enters both specifications, allowing separate identification under appropriate variation.

Second, if the utility specification δ_{ij} and the search cost specification μ_{ij} include excluded covariates, then the coefficients on covariates that enter only one of the two components can be identified without relying on additional assumptions or moment conditions. For example, in my specification, the distance between a consumer's residence and a pharmacy enters only the search cost equation and not the utility function.

Although μ_{ij} follows a parametric log-exp functional form (Equation 11), this choice imposes minimal constraint in practice. For moderate values (e.g., $\lambda > 3$), the approximation $\log(1 + \exp(\lambda_0)) \approx \lambda_0$ holds closely. For example, $\log(1 + \exp(3)) = 3.05$. In my estimation, the average estimated $\hat{\mu}_{ij}$ is 7.31. Thus, functional form is not the primary source of identification.

Identification does not require a closed-form reservation rule. The key is that any purchase at a non-default pharmacy reveals search: the buyer must have found the non-default's utility, net of the extra search cost, to beat both the default and the outside option. Because the default is always searched, its realization is known to the consumer (as is the outside option), whereas non-defaults are evaluated only after incurring an additional search cost. This asymmetry, between known default realization and non-default utility net of search, provides the variation needed to recover the search-cost component even without closed-form reservation values.

Correctly specifying the default option is essential. I define the default pharmacy as the one where the patient most frequently filled prescriptions prior to the index visit. For non-controlled substances, providers typically transmit prescriptions to the patient's designated pharmacy on file. A consistent fill history strongly suggests intentional selection. Misclassifying the default, for example by assuming it is simply the nearest pharmacy, would bias the estimates. If patients are more likely to fill prescriptions at their true default than at the nearest pharmacy, conditional on availability, this would overstate search intensity and understate true search costs.

VI. Results

This section proceeds in three steps. First, I present the estimated coefficients obtained from the estimation procedure and discuss their interpretation. Second, I show how the distribution of preference parameters and effective search costs differs between patients whose default pharmacy carries buprenorphine and those whose does not. Third, I simulate search behavior based on the estimated model. Finally, Appendix D discusses the internal validity of these results.

A. Estimated Coefficients

Table 1 reports the estimated coefficients from the preference and search components of the model. In the estimated preference parameters $(\hat{\delta})$, a higher value reflects greater utility derived from pharmacy and consuming buprenorphine i.e., a greater likelihood of filling the prescription. In contrast, higher values in the search cost parameters $(\hat{\mu})$ indicate greater barriers to search.

Price and Travel Distance: Patients are generally not price-sensitive, with a small but statistically significant negative coefficient on out-of-pocket price. Medicaid recipients, however, are substantially more price-sensitive, consistent with lower incomes in this population. Search costs increase with travel distance, both from the patient's residence and from their default pharmacy. This effect is more pronounced in urban areas, where each additional mile traveled from the residence or default pharmacy corresponds to a greater increase in disutility. Quantitatively, each additional mile of travel imposes an effective search cost of approximately \$0.62 from the patient's residence and \$0.32 from the default pharmacy. These values represent 2.9% and 1.6%

of the average out-of-pocket price of buprenorphine, respectively.²³ While these figures may appear modest, they imply far greater search burdens relative to price than in other markets. For comparison, Moraga-González, Sándor and Wildenbeest (2023) estimate that in the European car market, traveling one kilometer imposes an average cost of €148. Given an average vehicle price of €19,917, this amounts to only 1% of the product's value per mile traveled.

Provider Characteristics: Coefficients on provider characteristics, when entering through the utility function, should not be interpreted as prescribing behavior. Instead, they capture how provider attributes shape the perceived value of treatment. For example, patients assigned to female providers exhibit lower utility, potentially reflecting patient-side bias or differences in how treatment is communicated. In contrast, patients matched with more experienced providers or those specializing in addiction medicine report significantly higher utility, consistent with higher treatment quality or better alignment with patient needs. Provider characteristics also enter the search cost equation. One interpretation is that more experienced providers may help reduce patient uncertainty about buprenorphine availability. For example, they might direct patients to better-stocked pharmacies, thereby lowering effective search costs. However, the estimates do not support this mechanism: experience is not associated with lower search costs.

Health History: Health history variables reveal important heterogeneity in both preferences and search costs. Patients whose OUD diagnosis appears earlier in the coding hierarchy (i.e., more likely to be the primary diagnosis) exhibit higher utility from buprenorphine and face lower search costs. However, we observe what could be interpreted as adverse selection: patients with more severe health histories (i.e., higher past opioid use, mental illness, and substance use disorder) derive lower utility from treatment. This suggests that those who need buprenorphine most may perceive it as less valuable or be less responsive to treatment, posing a concern for equity and effectiveness. Interestingly, the relationship reverses in the search cost estimates: sicker patients face lower search costs. One possible explanation is that patients with multiple comorbidities needs fill multiple prescriptions and may already be familiar with pharmacy availability through other drugs, thus reducing marginal search costs. A greater number of diagnosed

 $^{^{23}}$ Estimates are scaled by the share of patients on Medicaid and Medicare, and distances are population-weighted based on urban ZIP code distributions.

Table 1— Estimated Coefficients

Preference Coefficients in δ			Search Coefficients in μ		
Variable	Estimate	SE	Variable	Estimate	SE
Out-of-Pocket Price	-0.005	2.7e-7	Distance from Residence	0.044	6.5e-5
\times Medicaid	-0.408	0.002	\times Urban	0.107	8.6e-5
\times Medicare	-0.006	5.0e-7	Distance from Default	0.039	4.8e-5
			\times Urban	0.039	6.4e-5
			Urban	0.036	0.027
Provider Characteristics					
Female Provider	-0.118	0.002	Female Provider	0.082	0.002
log(Experience)	0.043	4.0e-5	log(Experience)	0.045	4.0e-5
Addiction Medicine Specialist	0.025	0.004	Addiction Medicine Specialist	-0.006	0.004
\times log(Past Opioid Consumption)	-0.033	2.6e-4	\times log(Past Opioid Consumption)	-0.055	3.3e-4
\times log(Past Mental Illness)	-0.002	0.019	\times log(Past Mental Illness)	-0.004	0.022
\times log(Past Substance Use Disorder)	0.000	0.012	\times log(Past Substance Use Disorder)	-0.005	0.012
Health History					
OUD First Occurrence Position	-1.376	8.8e-4	OUD First Occurrence Position	0.469	0.003
# of All Conditions	0.026	1.2e-4	# of All Conditions	-0.061	1.2e-4
log(Past Opioid Consumption)	-0.314	5.4e-5	log(Past Opioid Consumption)	-0.242	7.4e-5
log(Past Mental Illness)	-0.035	0.004	log(Past Mental Illness)	-0.025	0.005
$\log({\it Past Substance Use Disorder})$	-0.038	0.002	$\log({\it Past Substance Use Disorder})$	-0.019	0.003
Demographics					
Black	-0.034	0.008	Black	0.026	0.007
Age	0.160	1.0e-4	Age	0.014	1.3e-4
Age Squared	-0.002	1.4e-8	Age Squared	-0.000	1.9e-8
Gender	-0.040	0.002	Gender	-0.045	0.002
Medicaid	-0.103	0.006	Medicaid	0.050	0.002
Medicare	-0.059	0.007	Medicare	0.032	0.008
Above High School	0.009	0.186	Above High School	0.026	0.328
$\log(\text{Income})$	0.126	0.010	$\log(\text{Income})$	0.313	0.008
Unemployment	0.005	0.878	Unemployment	-0.003	1.580
Poverty	0.013	0.168	Poverty	-0.005	0.160
Pharmacy Attributes					
Same Affiliation as Default	0.094	0.030	Same Affiliation as Default	-0.086	0.034
CVS	-0.041	0.056	CVS	0.035	0.060
Walgreens	-0.122	0.032	Walgreens	0.099	0.035
Rite Aid	0.169	0.029	Rite Aid	-0.151	0.031
Walmart	0.040	0.032	Walmart	-0.030	0.038
Albertson's	-0.030	0.031	Albertson's	0.044	0.033
Fred Meyer	-0.027	0.032	Fred Meyer	0.019	0.036
Bartell	-0.005	0.033	Bartell	0.020	0.036
Regional Chain	-0.010	0.037	Regional Chain	0.009	0.040
Constant	0.021	0.998	Constant	0.019	0.912

Note: Table 1 presents coefficient estimates obtained via maximum likelihood. Columns 1–3 display preference parameters, while columns 4–6 report search parameters. Heteroskedasticity-robust standard errors are computed using the Huber–White sandwich estimator.

conditions also predicts lower search frictions, potentially for similar reasons.

Demographics: Racial and socioeconomic disparities are evident. Black patients derive lower utility from buprenorphine and face higher search costs. Older patients have higher treatment preferences but also higher search costs, possibly due to mobility. Female patients show slightly lower preference for treatment, but also lower search costs. Medicaid and Medicare recipients derive lower utility and face higher travel-related barriers. Patients living in higher-income areas derive higher utility but encounter higher search costs, potentially due to lower pharmacy density or higher opportunity cost.

Pharmacy Attributes: Most chain brands (e.g., CVS, Walgreens) show no significant effects in either the preference or search cost parameters. This is partly because the effects are absorbed by an indicator for affiliation with the patient's default pharmacy (e.g., if the default pharmacy is CVS, all CVS locations receive a value of 1 for this variable). Patients filling prescriptions at pharmacies affiliated with their default location derive higher utility and face lower search costs, likely reflecting established relationships and greater familiarity. Rite Aid is an exception: patients derive higher utility and also face higher search costs. This may reflect a chain-wide policy to carry buprenorphine and serve OUD patients more consistently.

Overall, the results reveal significant heterogeneity in both utility and search costs. Disadvantaged populations, such as Black patients and those with more severe OUD histories, face a dual burden. They derive lower utility from buprenorphine and encounter higher barriers to access. This disparity is especially troubling given their elevated need for treatment. In Appendix F, I report the estimated distributions of μ and δ at the patient–pharmacy pair level.

B. Patient Search Behavior

Although the model is not estimated using direct observations of search behavior, it generates predicted search patterns based on the estimated parameters. For each patient, I simulate utility draws from a Type I Extreme Value distribution with location parameter $\hat{\delta}_{ij}$, and draw effective search costs from the distribution in Equation 7, using the estimated location parameter $\hat{\mu}_{ij}$. I repeat this process 100 times per patient. Using Equation 5, I compute the reservation value for each option. Patients are then assumed to follow the Weitzman rule: they visit pharmacies in

descending order of reservation values and stop once the realized utility exceeds the reservation value of the next-best option.

I focus on two key dimensions of search behavior: (i) the number of pharmacies visited and (ii) the incidence of *suboptimal searches*, defined as searches that result in visiting a pharmacy without buprenorphine. While such visits do reveal availability information, they ultimately do not expand the patient's choice set.

Figure 4a shows that patients exhibit limited search behavior. Over 60% of patients search only once, which by assumption corresponds to searching only their default pharmacy. This means the majority of patients do not actively engage in additional searching even when they encounter an unavailable pharmacy.²⁴ Such concentrated search behavior suggests substantial search costs that may significantly impede access to treatment. To quantify the magnitude of these search costs, I conduct a simple counterfactual analysis: eliminating search costs would increase the buprenorphine purchase rate by 38 percentage points, from 45% to 83%.

Figure 4b examines the extent to which patients visit pharmacies that do not have buprenorphine available. Patients are grouped by the total number of pharmacies visited, and for each group, the figure plots the average share of visits made to pharmacies without supply. For example, among patients who visited two pharmacies, one visit on average was to a pharmacy without buprenorphine, meaning that 50 % of their searches were unsuccessful. The share of unsuccessful visits increases with the total number of pharmacies visited. Among patients who visited only one pharmacy, 40 % encountered a location without supply. Among those who visited ten pharmacies, 78 % of visits were to locations that lacked buprenorphine. This pattern indicates that patients who search extensively are primarily seeking availability rather than engaging in search for match value, even though the model permits search for match value.

VII. Counterfactuals

Patients seeking buprenorphine face two distinct but reinforcing barriers: high physical search costs and substantial information frictions. I evaluate a series of counterfactual policies designed

²⁴The predicted distribution of search frequency is consistent with prior findings in other high-stakes settings. For example, Ambokar and Samaee (2019) show that in the mortgage refinancing market, 59% of U.S. borrowers make only one inquiry, and about 80% make fewer than two. This degree of inaction is comparable to the behavior observed in this context.




Figure 4. Predicted Consumer Search Behaviors

Note: This figure presents predicted patient search behavior based on the estimated model. Panel (a) shows the distribution of the number of pharmacies visited. Panel (b) reports the average share of visits made to pharmacies without buprenorphine available, conditional on the total number of pharmacies visited. Search behavior is simulated using estimated utility and search cost parameters. For each simulated patient, realized utilities and reservation values are drawn, and the Weitzman rule is applied to determine the sequence and number of searches.

to reduce these barriers, each aimed at improving initiation at a fraction of the cost of mandating universal availability.

The first set of counterfactuals considers pure information provision. Systems that allow prescribers to observe real-time pharmacy inventory or centralized platforms that report availability would reduce uncertainty without altering the underlying search process. By eliminating guesswork at the start of the search, such policies can substantially increase initiation.

A second policy focuses on electronic prescribing. During the study period, controlled substance prescriptions were typically handwritten and could not be transmitted in advance. E-prescribing allows physicians to send prescriptions directly to pharmacies, enabling availability to be confirmed before patients begin their search. At the same time, however, e-prescriptions cannot be transferred without contacting the prescriber for a new one. This rigidity increases the cost of switching if the first pharmacy is out of stock. The policy therefore reduces uncertainty but simultaneously raises physical search costs.

Finally, I consider how searching costs under mandatory e-prescribing might be offset through direct financial incentives to patients. For instance, providing small monetary transfers to cover travel or time costs would encourage patients to continue searching when their default pharmacy is unavailable. By lowering the effective cost of additional search, such incentives complement the informational benefits of e-prescribing and further improve treatment initiation.

A. Information Provision Counterfactuals

I examine three counterfactual policies that provide certainty about pharmacy availability:

- 1) Universal Availability Mandate. This benchmark scenario mandate that all pharmacies are required to carry buprenorphine. By eliminating uncertainty, this policy minimizes effective search costs. In addition, it expands the choice set as all pharmacies carry the drug. While useful as an upper bound on policy effectiveness, its implementation may be infeasible due to operational and financial constraints.
- 2) **Perfect Information via Providers.** In this scenario, the state maintains a real-time database of pharmacy-level availability that prescribers can access and share with patients

at the point of care. This approach preserves pharmacy discretion, as inventory is not advertised directly to patients. The intervention reduces effective search costs by eliminating uncertainty but does not increase the number of pharmacies stocking buprenorphine. Patients only search pharmacies that carry buprenorphine.

3) Partial Information via Informal Networks. Even in the absence of a centralized inventory system, prescribers may form informal agreements with a small number of local pharmacies, e.g., the five closest, to reliably carry buprenorphine for their patients. In this scenario, the prescriber shares availability information only for these pharmacies. Effective search costs fall for the subset of pharmacies in the network but remain unchanged elsewhere. Patients are free to visit any pharmacy, but they know with certainty that at least some pharmacies will carry buprenorphine.

All three policies reduce information frictions by increasing certainty about which pharmacies carry buprenorphine. Under a universal availability mandate, patients believe that every pharmacy stocks the medication. In the informal network scenario, patients are certain that a subset of nearby pharmacies—e.g., the five closest—will have buprenorphine available. These interventions shift patients' beliefs about availability, altering the belief parameters ϕ_{ij} .

I assume that patients form rational expectations about local availability to circumvent the problem that beliefs are not separately identified in the model. Let $\hat{\phi}_i$ denote patient *i*'s perceived probability that a randomly selected pharmacy within their choice set has buprenorphine available. For instance, if ten pharmacies are within reach and three carry buprenorphine, then $\hat{\phi}_i = 0.3$. Figure E1 reports the cross-patient distribution of $\hat{\phi}_i$. Under this framework, the effective search cost is scaled by availability beliefs: $\tilde{c}_{ij}/\hat{\phi}_i$, where \tilde{c}_{ij} is the baseline effective search cost for pharmacy j.

To simulate search, purchase behavior and consumer surplus, I draw 100 realizations from the adjusted search cost distribution $F_{\tilde{c}}(\tilde{c}_{ij}/\hat{\phi}_i)$ for each patient. Given these simulated search costs and the estimated utility component $\hat{\delta}_{ij}$, I compute the reservation value using Equation 5. The latent utility w_{ij} is then calculated as the minimum of the reservation value and the realized utility, where the latter is drawn from a Type I Extreme Value distribution with the location paramter, $\hat{\delta}_{ij}$. For purchase behavior, patients select the pharmacy with the highest w_{ij} . For




Figure 5. Buprenorphine Purchase Counterfactuals

Note: This figure is generated by simulating each patients' w_{ij} values by simulating search cost draw from the distribution of $F_{\tilde{c}}(\tilde{c}_{ij}/\hat{\phi}_i)$. Then to obtain the population share purchasing buprenorphine under three counterfactual scenarios: (i) all pharmacies offer buprenorphine, (ii) doctors have full knowledge of pharmacy availability, and (iii) doctors know five nearby pharmacies that offer the medication. The baseline reflects actual purchase shares observed in the data. The number on the top of each bar indicates percentage increase compare to the baseline scenario under that scenario.

search behavior, I apply the Weitzman rule: patients visit pharmacies in descending order of reservation values, stopping once the realized utility at a visited pharmacy exceeds the reservation value of the next-best alternative. Consumer surplus is measured as the realized utility from the chosen product minus the cumulative costs of all pharmacies visited, scaled by the price coefficient to express the result in dollar terms.

Figure 5 displays the predicted share of patients who successfully obtain buprenorphine under each policy scenario. If all pharmacies offered buprenorphine, the purchase rate would rise to 58%. Mandating universal availability raises purchase rates by 29% relative to baseline, but such a policy may be financially or logistically infeasible. From a state policy perspective, a more practical alternative is to equip prescribers with real-time information on pharmacy availability and inform patients of nearby pharmacies that carry buprenorphine at the point of care. Unlike




Figure 6. Number of Searches Counterfactuals

Note: This figure is simulated similarly to Figure 5, but it shows the predicted number of pharmacies searched by patients under each policy scenario. The baseline reflects actual search behavior observed in the data.

a universal availability mandate, this policy does not expand the choice set but substantially reduces search costs by removing uncertainty. It increases the predicted purchase share by 17%, about 58% of the effect of universal availability. The key distinction is that universal availability ensures every patient has a default pharmacy that stocks buprenorphine, whereas the information policy improves initiation solely by eliminating uncertainty. This highlights that uncertainty accounts for most of the gains in treatment initiation.

Even in the absence of system-wide coordination, smaller-scale interventions can yield substantial benefits. If providers maintain informal relationships with five pharmacies known to dispense buprenorphine and communicate this information to patients, treatment initiation increases by 7%. Although patients may still face uncertainty at other locations, the presence of a few known options lowers search costs and improves access. This intervention can be implemented at the individual provider level, offering a low-cost, decentralized alternative to broad mandates. A similar pattern emerges for consumer surplus, as shown in Figure E2.

Figure 6 illustrates how patient search behavior responds under different policy scenarios. First, mandating universal availability increases search activity by 12%. This policy lowers effective search costs and expands the choice set, but the increase remains modest. Patients who find a pharmacy with buprenorphine may still continue searching because additional visits are now less burdensome due to reduced effective search costs. Yet most searches remain driven by availability rather than by match value. Second, when providers inform patients about availability at all pharmacies, searches fall by 37%. Patients avoid pharmacies without supply, particularly their default pharmacy if it does not carry buprenorphine, where previously all patients visited by default. Lastly, the five-pharmacy information policy generates a moderate 13% increase in searches, comparable to universal availability. This reflects the same group of patients who continue searching beyond the first available option because reduced effective search costs make additional visits to the five informed pharmacies less burdensome. As before, most patients remain motivated by availability rather than by match value.

A key assumption in the policy analysis is that patients form rational expectations about local buprenorphine availability. This may be strong: if patients are naïve and presume pharmacies always carry the medication, then observed behavior would reflect only physical search costs, and information provision would appear ineffective. To assess robustness, I conduct a sensitivity analysis that rescales patients' perceived availability. Specifically, I vary the belief parameter $s\hat{\phi}_i$ with $s \in [0.1, 2]$, where s < 1 corresponds to pessimistic beliefs and s > 1 to optimistic beliefs, and recompute initiation under the "prescriber-informed availability" counterfactual. Results (Figure E3) show two patterns. First, the response is asymmetric: moving from baseline toward pessimism raises initiation more than an equally sized move toward optimism reduces it. Second, even under very optimistic beliefs (large s), information provision still raises initiation by a modest amount.

As an additional check, I estimate ϕ_{ij} using the realized pharmacy stocking history that patients would rely on if they formed rational beliefs about the probability that a pharmacy carries buprenorphine. This exercise treats patients as if they had perfect information about past availability when forming their beliefs. The resulting estimates closely align with the baseline

assumption of rational expectations regarding local buprenorphine availability.²⁵ In Appendix G, I report all counterfactual results (including the next two analyses) under this alternative assumption.

B. Electronic Prescribing

By 2025, all U.S. states will require electronic prescriptions for controlled substances, including buprenorphine. This policy aims to reduce diversion and improve traceability by linking each prescription to a specific patient, provider, and pharmacy. In doing so, it addresses one of the early drivers of the opioid crisis, unregulated prescribing and dispensing by rogue actors.

A key benefit of e-prescribing is that it eliminates ex-ante uncertainty about pharmacy availability. Under the paper-based system, patients needed to visit pharmacies to determine whether buprenorphine was available. With e-prescriptions, orders are transmitted directly to the selected pharmacy, which then confirms whether it can be filled. This reduces information uncertainty and thus lowers the overall effective search cost faced by patients.

However, e-prescribing also introduces a barrier to access. For controlled substances, e-prescriptions generally cannot be transferred between pharmacies without the prescriber issuing a new prescription.²⁶ If the selected pharmacy is out of stock, the patient must re-contact the prescriber to obtain a new prescription, increasing search costs. This burden could be consequential for buprenorphine, given that many pharmacies do not carry it.

To quantify this trade-off, I model e-prescription as a combination of two opposing forces: (i) a reduction in uncertainty and (ii) an increase in search costs if the prescription must be redirected. Although the exact burden of re-contacting a provider is unknown, I assess how much search costs would need to increase to offset the informational benefit of certainty. Let s denote the multiplier applied to the search cost location parameter μ . For example, s=2 implies that costs are doubled. As before, uncertainty is captured by the average perceived availability $\hat{\phi}_i$. Following the previous procedure, I simulate the search costs by drawing from the distribution

²⁵The "universal availability" counterfactual is highly sensitive to beliefs. Many pharmacies historically did not carry buprenorphine, so patients plausibly assign them very low availability. Under the policy, perceived availability at those locations jumps from near zero to one, eliminating the uncertainty component of search and sharply lowering effective search costs. This is specific to the universal-availability scenario; in the other counterfactuals, pharmacies with low prior availability remain unavailable and therefore do not enter the choice set, as they do not carry buprenorphine.

 $^{^{26}}$ This applies only to controlled substances; it does not apply to other medications.




Figure 7. Trade-Offs Under E-Prescription

Note: This figure illustrates how outcomes change when patients face higher average search costs. Specifically, it shows the level of additional search costs at which the e-prescription policy performs worse than the baseline. The baseline reflects actual purchase shares observed in the data. Additional search costs are incorporated by scaling the location parameters of the search cost distribution, μ_{ij} , such that the effective search cost becomes $s\mu_{ij}$, where s is the scaling factor. After doing this, it draws from $F_{\bar{c}}(\tilde{c}_{ij}/\hat{\phi}_i)$ to simulate the search costs then formulate the reservation values and realized utility.

 $F_{\tilde{c}}(\tilde{c}_{ii}/\hat{\phi}_i)$. The reservation values and realized utility are then constructed accordingly.

Figure 7 shows how predicted buprenorphine initiation share vary with the search cost multiplier s. The baseline reflects observed initiation level under the paper-based regime. When s=1, the predicted initiation matches the "perfect information via providers" scenario discussed previously. As s increases, initiation declines, reflecting the additional cost of coordinating with providers to redirect prescriptions when a pharmacy lacks inventory.

Importantly, for the costs of e-prescription to outweigh its informational benefits, meaning for predicted initiation share to fall below the baseline, the search cost would need to be doubled. This magnitude is unlikely in practice: even accounting for delays and provider inaccessibility, the effort of contacting a prescriber to reissue a prescription is unlikely to be twice as burdensome as physically visiting a pharmacy. Overall, the e-prescribing mandate is therefore likely to

raise buprenorphine initiation by reducing the informational frictions that suppress treatment initiation under the paper-based system.

C. Monetary Incentives

The preceding counterfactuals show that reducing informational frictions can substantially increase buprenorphine initiation. Yet even eliminating uncertainty does not close the gap between initiation for OUD and the initiation share observed for other chronic conditions. The remaining shortfall is likely due to high baseline search costs, such as the burden of traveling even one additional mile. I therefore ask how much direct monetary incentives to patients could further increase initiation by offsetting these physical search costs.

For context, several states, including California and West Virginia, have begun piloting programs that provide modest financial incentives to promote treatment adherence.²⁷ These programs typically provide gift cards of \$10–\$20 at treatment milestones, subject to a federal cap of \$599 per patient per year. Inspired by these initiatives, I evaluate whether small subsidies could be effective at the very first step of treatment, ensuring that patients fill their initial buprenorphine prescription. Because the first fill substantially increases the likelihood of continued adherence, targeting this margin may offer a cost-effective lever.

I focus on Medicaid patients (same as in the mentioned states), for whom such an intervention is likely to be most effective. They represent roughly 40% of all individuals with OUD, are more price-sensitive than other groups in the estimated model, and remain the least likely to initiate buprenorphine. To model the counterfactual, I augment Medicaid patients' utility with a direct subsidy: I add x, scaled by the Medicaid-specific price coefficient, to the x term in Equation 10. Purchase probabilities are then recomputed using Equation 12. Because e-prescription has since become mandatory, it is important to examine how much monetary incentive is required in today's environment with e-prescription in place. Following the earlier approach of lowering the uncertainty about pharmacy availability, I then calibrate the additional burden from e-prescribing as a 20% increase in the search cost location parameter μ .

²⁷California's Recovery Incentives Program, implemented under the CalAIM Section 1115 Demonstration, provides contingency management incentives in the form of gift cards up to \$599 annually for patients with substance use disorder (California's Contingency Management Benefit). West Virginia's Section 1115 SUD Waiver similarly authorizes the use of gift card incentives for substance use disorder treatment under Medicaid Section 1115 Demonstration.




Figure 8. Medicaid Rebate Counterfactual

Note: The figure plots the purchase probability for Medicaid patients under varying subsidy levels, with and without e-prescription. Subsidies are modeled as direct utility transfers equivalent to a reduction in out-of-pocket cost. For the scenario without e-prescription, probabilities are calculated using the estimated parameters from Equation 12. For the e-prescription scenario, effective search costs are calibrated to be 20% higher to reflect the need for prescriber reauthorization when a pharmacy is out of stock. Purchase decisions follow the Weitzman search rule under simulated search costs $F_{\tilde{c}}(\tilde{c}_{ij}/\hat{\phi}_i)$.

Figure 8 reports share of people initiating buprenorphine for Medicaid patients across subsidy levels, with and without e-prescription. Two results stand out. First, in the baseline without any information effect (i.e., no e-prescription), a \$20.24 subsidy raises the initiation share to 70%. 28 Second, once e-prescription is in place, a subsidy of about \$13.49 is sufficient to lift initiation to levels comparable to other chronic conditions. Both amounts fall within the range of existing gift-card incentives, indicating the policy is feasible and scalable. Given the nationwide shift to e-prescribing, these results suggest that modest, Medicaid-targeted subsidies are a practical tool to close the "last mile" to treatment. The paper's policy recommendation is therefore to pair the existing e-prescribing infrastructure with small, first-fill gift cards to substantially increase buprenorphine initiation among Medicaid patients.

²⁸This is noteworthy because even under perfect naiveté, where patients believe every pharmacy is available and thus there is no information mechanism, \$20 still suffices.

VIII. Conclusion

Buprenorphine remains underused at treatment initiation because the "last mile" is difficult to navigate. Even with a prescription in hand, patients face two search frictions: the physical cost of visiting a pharmacy and uncertainty about which pharmacies actually stock the medication. To quantify the effective search cost, I develop a sequential search framework that recovers both preferences and search frictions using a novel identification strategy. The key assumption is that patients always search their default pharmacy, which normalizes its search cost to zero. A fill at a non-default pharmacy therefore reveals that its utility net of search cost exceeds that of the default, providing information on otherwise unobserved search. These deviations supply the variation needed to separately identify preferences and search frictions without observing search sequences.

The counterfactuals show that administratively simple tools can raise initiation to levels comparable to those for other chronic conditions. Because e-prescribing is now mandatory, its informational benefits are already in place; pairing those benefits with modest, targeted first-fill gift cards further reduces the physical search costs that deter initiation. These levers are policy-ready, low-cost, and scalable, and together can convert more prescriptions into treatment starts without new technologies or substantial fiscal commitments. In turn, they can expand access to life-saving treatment and improve continuation for patients with OUD.

Data Construction

A1. Patient Identification

The objective of this section is to identify patients who initiated buprenorphine treatment for OUD. Ideally, a patient's initiation is observed as a medical claim indicating an outpatient visit for OUD, followed by a corresponding pharmacy claim for a buprenorphine prescription. However, two challenges complicate this process: (i) buprenorphine can be prescribed without an explicit OUD diagnosis, and (ii) OUD diagnoses may lead to alternative treatments such as methadone, without generating buprenorphine pharmacy claims.

To address these issues, I apply the following data processing strategy. First, I exclude all patients with any evidence of prior OUD-related care (identified via ICD-9/10 diagnosis codes in medical claims) or buprenorphine prescriptions (identified via NDC codes specific to medications for opioid use disorder in pharmacy claims) before April 1, 2014. I focus exclusively on outpatient visits, removing encounters in emergency or institutional settings, and restrict to medical claims involving providers who prescribed buprenorphine during the same calendar year. This ensures the sample includes only new episodes of care likely associated with buprenorphine initiation.

Among these patients, I classify the first outpatient medical claim with an OUD diagnosis as a potential treatment initiation. If this visit is followed by a buprenorphine prescription fill within 14 days, I classify the patient as plausibly matched.

For patients who have a buprenorphine prescription without a preceding medical claim meeting the above criteria, I search for the closest outpatient visit within 14 days of the fill date. Preference is given to visits involving the same prescriber as recorded in the pharmacy claim, and among those, to the visit closest in time. This step helps address cases where the medical claim lacks an explicit OUD diagnosis, but the prescriber is known to have issued buprenorphine.

Finally, for patients with a qualifying OUD-related medical visit but no timely pharmacy match, I check for sustained follow-up with the same provider. If the patient returned to the same provider at least three times within the following 90 days, I infer that the care episode likely reflects a different treatment plan (e.g., methadone or non-pharmacologic intervention), and exclude these cases to avoid ambiguity.

This process yields three mutually exclusive groups of patients:

- 1) Matched OUD initiators: Patients whose first OUD-related medical visit is followed by a buprenorphine fill within 14 days. (22% of the first time patients)
- 2) Pharmacy-first initiators: Patients with a buprenorphine prescription who are matched to a proximate outpatient visit with the same prescriber, even if the medical claim lacks an OUD diagnosis. (22% of the first time patients)
- 3) Unmatched medical visits: Patients with OUD-related outpatient visits who do not initiate buprenorphine and do not show sustained provider follow-up suggestive of alternative treatment. (56% of the first time patients)

A2. Inventory Construction

To make the problem tractable, I adopt several simplifying assumptions when constructing pharmacy-level buprenorphine inventory. First, I aggregate across formulations and strengths. Buprenorphine is commonly dispensed either as a monotherapy or in combination with naloxone, most often in 2mg, 8mg, or 12mg doses. While substitution between brand-name and generic products is generally straightforward, clinical substitution between formulations, such as switching from buprenorphine alone to buprenorphine/naloxone, requires medical judgment and may not always be appropriate. Despite these differences, I aggregate all formulations into total milligrams of buprenorphine, treating them as interchangeable. This abstraction ignores strength-specific constraints. For example, pharmacies may not split high-dose tablets to fulfill low-dose prescriptions, and the approach likely overstates the availability of buprenorphine.

Second, the shipment data from ARCOS and the sales data from WA-APCD are not directly comparable. ARCOS captures all shipments to pharmacies, whereas the WA-APCD includes only pharmacy transactions for Washington residents covered by insurance, excluding both cash payments and out-of-state patients. On average, WA-APCD sales account for approximately 40% of total shipments. To assess whether a pharmacy has inventory available for a given patient, I evaluate the pharmacy's stock on the relevant date: the pharmacy claim date if available, or otherwise the medical claim date. For each pharmacy within the patient's choice

set, I calculate the difference between cumulative shipments over the prior seven days (scaled by 0.403 to reflect the share of sales captured in WA-APCD) and cumulative WA-APCD sales over the same window. A pharmacy is considered to have sufficient stock if the resulting estimate exceeds the patient's inferred need on that date. When a pharmacy claim is observed, I use the actual strength and duration of that claim to determine the patient's need. When no pharmacy claim is available, I impute demand as a 14-day supply of 8mg buprenorphine/naloxone, the most common treatment pattern observed in the data.




Figure B1. Number of Days Between Doctor Visits and Prescription Filled

Note: This figure displays the distribution of days between the first OUD-related medical visit and the first buprenorphine prescription fill, and between the first hypothyroidism-related visit and the first levothyroxine prescription fill. Patients medical conditions are sourced from WA-APCD.

ADDITIONAL STYLIZED FACTS

B1. Time to Fill a Prescription After Diagnosis

Access to treatment depends not only on whether a prescription is written, but also on how quickly patients are able to fill it. For some medications, patients may wait before filling a prescription, particularly if the drug is not urgently needed. For buprenorphine, however, the timing of initiation is critical. Patients must begin treatment at the onset of withdrawal, and delays can increase the likelihood of relapse. As a result, even short search frictions can carry meaningful clinical consequences.

To examine delays in access, I compare the time between diagnosis and prescription fill for two drugs: buprenorphine and levothyroxine. I restrict attention to first-time users of each medication. The timing measure is defined as the number of days between the first diagnosisrelated medical visit and the first fill of the corresponding prescription.

Figure B1 presents the distribution of fill times. Most patients in both groups fill their prescriptions within a few days. Among OUD patients, nearly 70% obtain buprenorphine within one day of diagnosis, compared to 55% of hypothyroidism patients who fill levothyroxine within the same window. The more rapid initiation of buprenorphine reflects its urgency in managing withdrawal, but also underscores the importance of pharmacy availability. When the default pharmacy does not carry buprenorphine, even brief delays in search may disrupt treatment initiation and risk relapse. To ensure that these differences are not driven by underlying differences in patient populations, I condition on prior opioid exposure: both samples are restricted to individuals with at least four prior opioid prescriptions. Results are similar when further conditioning on patients who eventually fill a buprenorphine prescription, though the sample size is smaller.

B2. Search Patterns Following Unavailability at the Default Pharmacy

This section presents stylized facts consistent with patients engaging in additional search when buprenorphine is unavailable at their default pharmacy. While I do not observe the search process directly, I examine subsequent purchase behavior to infer patterns.

I focus on patients who ultimately filled a prescription after encountering unavailability at their default pharmacy. For each such case, I calculate two distances for the pharmacy where the prescription was filled: the distance to the patient's residential ZIP code and the distance to the default pharmacy. If patients always check availability at the default pharmacy first, then the observed purchase can be interpreted as the result of continued search. Proximity to home reflects baseline preference for convenience, while proximity to the default pharmacy captures behavior consistent with search initiated after discovering unavailability.

Figure B2 shows the distribution of these two distances. Both are skewed toward shorter values, indicating that most patients fill prescriptions close to home and close to the default pharmacy. However, the distribution is more concentrated near zero for the distance to the default pharmacy. This suggests that while home proximity remains important, the location




Figure B2. Distance to Default and Residential ZIP Code for Purchased Pharmacy

Note: This figure shows the distribution of distances from the purchased pharmacy to the default pharmacy and the patient's residential ZIP code, conditional on the default pharmacy being unavailable. The sample includes patients who ultimately filled a buprenorphine prescription elsewhere.

of the default pharmacy also shapes the search path. This implies that patients continue their search from that point, rather than starting from home alone. These patterns justify a search-based model, as distance plays a clear and significant role in consumer choice.

B3. Health Outcomes

The paper's focus on increasing buprenorphine initiation share, but to what extend taking buprenorphine would improve the health outcomes? It is not clear to the literature, evidence largely builds on observational study and clinical evidences. However, the paper's has a unique design that allows to estimate the effect of buprenorphine treatment on health outcomes. The

problem with directly estimating patients with OUD that whether take or not take buprenorphine is that subject strongly to selection bias. Patients who take buprenorphine are likely to be more health-conscious, and thus may achieve the same outcomes without medication. To address this, I can use an instrument for this endogenous treatment, in that, patients are more likely to take buprenorphine if the default pharmacy has buprenorphine available, but the availability itself should not directly affect health outcomes except through treatment initiation. Further, together I can use a difference-in-differences (DiD) design to controls for fixed unobservables, which led to a more credible identification.

First, like the whole paper I focus on the first-time OUD patients with buprenorphine prescription. The treatment group consists of patients who receive buprenorphine, while the control group includes only those who never received buprenorphine after the buprenorphine prescription.²⁹ DiD specification as following,

(B1)
$$y_{it} = \beta \text{Buprenorphine Treatment}_{it} + \alpha_i + \gamma_t + \text{Doctor}_i + \text{ZIP}_i + \epsilon_{it}$$


where y_{it} represents the number of emergency room visits per quarter, Buprenorphine Treatment_{it} is a binary indicator for treatment, and α_i and γ_t are individual and time fixed effects. I further control for doctor and ZIP code fixed effects.

To address endogeneity concerns, I instrument for buprenorphine treatment using its availability at the default pharmacy. Patients are more likely to initiate treatment when buprenorphine is available at their default pharmacy, but availability itself should not directly affect health outcomes except through treatment initiation. The instrumental variable regression is:

(B2) Buprenorphine Treatment_{it} =
$$\pi$$
Default Availability_{it} + $\tilde{\alpha}_i + \tilde{\gamma}_t + \tilde{\text{Doctor}}_i + \tilde{\text{ZIP}}_i + \eta_{it}$

Figure B3 presents ER visit trends for treatment and control groups. The raw trends (Figure B3a) reveal several key patterns. First, patients do not immediately receive OUD treatment

²⁹Later-treated patients are excluded from the control group to prevent bias from differential health trends.



Note: This figure plots trends in emergency room visits for patients who received a buprenorphine prescription and either initiated treatment (treatment group) or did not (control group). ER visits are identified in the WA-APCD medical claims using place-of-service codes, with multiple claims on the same date counted as a single visit. Panel (a) shows raw trends, while Panel (b) restricts to patients between the 60th and 90th percentiles of pre-treatment ER visit frequency.

Figure B3. Emergency Room Visits Per Quarter Trends by Treatment Groups

after an ER visit; there is a buildup process before they enter care. Second, receiving an OUD diagnosis is the first step toward recovery, and ER visits decline sharply after diagnosis, regardless of whether the patient starts buprenorphine treatment. Finally, patients who eventually receive buprenorphine already have lower ER visit rates before treatment. This last point has important implications. If the full sample are used directly, the estimated effect of buprenorphine may be underestimated, particularly when reductions in ER visits are more pronounced for those starting from a high baseline.³⁰

To address this concern, I ensure that patients and patients are comparable between treatment and control groups. For example, if I only focus on those between the 60th and 90th percentiles of pre-treatment ER visit frequency, as shown in Figure B3b, I can ensure that treatment and control groups are more similar at baseline. In light of these, I stratify patients into quintiles based on pre-treatment ER visits to ensure that treatment and control groups are comparable, mitigating concerns about the initial baseline.

Figure B4 presents estimated treatment effects from DiD and instrumented difference-indifferences (DiD-IV) regressions. Across all quintiles, most estimates show statistically significant reductions in emergency room visits. In many cases, these declines exceed 20% relative to the pre-treatment mean for the treated group.

The DiD-IV estimates are consistently larger in magnitude, though not statistically different from the DiD estimates. The largest differences arise in the bottom two quintiles, where DiD estimates are close to zero while DiD-IV identifies substantial reductions in emergency room visits. This reflects the fact that DiD-IV scales the DiD estimate by the first-stage effect of the instrument (i.e., divides by β in Equation B2).

The DiD-IV specification captures an average causal response among compliers, defined as patients whose treatment initiation depends on whether buprenorphine is available at their default pharmacy. It yields a properly weighted average of treatment effects for this subgroup. In contrast, the standard DiD estimate reflects a mixture of compliers and always-takers, the latter being patients who would obtain buprenorphine regardless of pharmacy availability. Because

³⁰A useful analogy is weight loss treatments: Suppose drug A is given to severely overweight individuals, while drug B is taken by those who are only moderately overweight. If weight loss is greater for those taking drug A, this does not necessarily mean it is the more effective treatment. It could simply reflect differences in initial weight.




Figure B4. Taking Buprenorphine Effects on Emergency Room Visits

Note: This figure plots the estimated effects of buprenorphine use on emergency room visits across quintiles of pre-treatment ER utilization. The DiD estimates are obtained from Equation B1, while the DiD-IV estimates come from an instrumental variables specification that uses buprenorphine availability as an instrument for treatment and substitutes the predicted values into the DiD framework. Ninety-five percent confidence intervals are clustered at the patient level.

always-takers are likely more motivated to seek treatment, their average treatment effect is smaller. As a result, the DiD estimate is smaller than the DiD-IV estimate.

Proof

Theorem 1—Eventual Purchase: Let $w_j \equiv \min\{\mathbf{1}\{j \in A\}u_j + (1 - \mathbf{1}\{j \in A\})(u_0 - \iota), r_j\}$ for each j. Given u_0 , the consumer purchases product j if and only if $w_j > u_0$ and $w_j > w_k$ for all $k \neq j$.

Proof: Sufficiency: If $w_j > u_0$, product j is available (if unavailable, $w_j < u_0$). Given product j is available, she purchases a product j because she is willing to visit at least one seller $(r_j > u_0)$ and make a purchase $(u_j > u_0)$. It remains to show that if $w_j > w_k$ when product j is available, then product k is not chosen.

- Case 1: Suppose product k is unavailable $(\mathbf{1}\{k \in A\} = 0 \implies w_k < u_0)$, she has no incentive to visit seller k or purchase product k. I only need to consider when product k is available.
- Case 2: Suppose $r_k \leq u_k$, then $w_k = r_k^*$. The consumer visits seller k only after seller j because $r_k \geq w_j > r_k$. However, once she visits seller j, she has no incentive to visit seller k because $u_j > r_k$.
- Case 3: Suppose $r_k > u_j$, which implies that $w_k = z_k$. In this case, even if she visits seller j, she either recalls a previous product $(u_j > z_k)$ or continues to search $(r_j > z_k)$ and finds a better product $(u_j > z_k)$.

Necessity: If $w_j < u_0$, then seller j is neither visited $(r_j < u_0)$ nor product j is purchased $(\mathbf{1}\{j \in A\}u_j < u_0)$, regardless of availability.

If $w_j < w_k$ for some $k \neq j$, and product j is unavailable, product j cannot be chosen $(w_j < u_0)$. If j is available, $w_j < w_k$ implies product k is also available as $0 < w_j < w_k$. By the same logic as above (case 2 and case 3), the consumer does not purchase product j.

INTERNAL VALIDATION

Figure D1 compares predicted and observed buprenorphine purchases, conditional on whether the default pharmacy stocks the medication, across three categories: purchase at the default pharmacy, purchase at a non-default pharmacy, and no purchase. The largest discrepancy, no




Figure D1. Internal Validation

Note: Figure D1 compares the observed and predicted number of patients who purchase at their default pharmacy, at a non-default pharmacy, or do not purchase at all, under two scenarios: when buprenorphine is available at the default pharmacy and when it is not. Predicted purchases are obtained by aggregating individual purchase probabilities, calculated using the estimated coefficients from Table 1 in Equation 12, to the pharmacy level.

more than 3%, arises when buprenorphine is unavailable at the default pharmacy but purchased at a non-default pharmacy. Because these categories are not moments directly targeted in estimation, the close alignment between predictions and observed outcomes provides a strong validation of the model and supports its use in generating counterfactuals.

Additional Tables and Figures

Table E1— Summary Statistics

Variable	Mean	St. Dev.	\mathbf{Min}	P25	Median	P75	Max
Individual Variables							
OUD First Occurrence Position	2.37	2.26	1.00	1.00	1.00	3.00	15.00
Past Opioid Consumption (MGE)	14,885.33	45,679.43	0.00	0.00	0.00	3,300.00	1,991,675.00
Past Mental Illness Diagnoses	0.70	2.46	0.00	0.00	0.00	0.00	92.00
Past Substance Use Disorder Diagnoses	0.93	3.41	0.00	0.00	0.00	1.00	92.00
# of All Conditions	4.05	2.25	1.00	3.00	4.00	5.00	25.00
Black	0.06	0.23	0.00	0.00	0.00	0.00	1.00
Age	41.85	14.57	0.00	30.00	39.00	53.00	90.00
Gender	0.52	0.50	0.00	0.00	1.00	1.00	1.00
Medicaid	0.55	0.50	0.00	0.00	1.00	1.00	1.00
Medicare	0.14	0.34	0.00	0.00	0.00	0.00	1.00
Above High School	0.90	0.06	0.00	0.88	0.91	0.94	1.00
Income	35,167.69	12,107.95	3,594.00	27,798.00	32,703.00	38,942.00	189,150.00
Unemployment	0.06	0.02	0.00	0.04	0.06	0.07	1.00
Poverty	0.13	0.06	0.00	0.08	0.12	0.16	1.00
Female Provider	0.39	0.49	0.00	0.00	0.00	1.00	1.00
Provider Experience	8.84	3.28	0.00	7.00	10.00	11.00	14.00
Addiction Medicine Specialist	0.19	0.39	0.00	0.00	0.00	0.00	1.00
Urban	0.92	0.28	0.00	1.00	1.00	1.00	1.00
Pharmacy Variables							
CVS	0.03	0.17	0.00	0.00	0.00	0.00	1.00
Walgreens	0.09	0.29	0.00	0.00	0.00	0.00	1.00
Rite Aid	0.09	0.29	0.00	0.00	0.00	0.00	1.00
Walmart	0.05	0.21	0.00	0.00	0.00	0.00	1.00
Albertsons	0.17	0.37	0.00	0.00	0.00	0.00	1.00
Fred Meyer	0.06	0.24	0.00	0.00	0.00	0.00	1.00
Bartell	0.05	0.22	0.00	0.00	0.00	0.00	1.00
Regional Chain	0.06	0.24	0.00	0.00	0.00	0.00	1.00
Individual-Pharmacy Variables							
Price	21.33	78.19	0.00	0.00	8.25	8.30	595.80
Distance from Default	15.52	9.43	0.00	8.42	14.52	21.16	50.00
Distance from Residence	15.07	8.26	0.02	8.90	14.78	20.48	50.00

Note: This table presents summary statistics for the main variables used in the analysis. The variables are grouped into individual characteristics, pharmacy characteristics, and individual-pharmacy level variables. The mean, standard deviation, minimum, 25th percentile (P25), median, 75th percentile (P75), and maximum values are reported for each variable.




Figure E1. Distribution of Estimated Belief Parameter $\hat{\phi}$

Note: This figure plots the distribution of the estimated belief parameter $\hat{\phi}_i$, patient i's belief that pharmacy j carries buprenorphine, at the patient level. Distributions are shown separately for patients whose default pharmacy stocks buprenorphine and for those whose default does not. Beliefs are formed under a rational-expectations assumption using neighborhood availability: the prior for each pharmacy equals the share of nearby pharmacies that stock buprenorphine (e.g., if 3 of 10 local pharmacies stock it, the prior at any given pharmacy is 0.30).




Figure E2. Consumer Surplus Counterfactuals

Note: This figure is simulated in a manner similar to Figure 5. Consumer surplus is calculated as the difference between realized utility and cumulative search costs, scaled by the price coefficient.




Figure E3. Pessimistic vs. Optimistic Beliefs

Note: The figure plots the predicted buprenorphine purchase share as the belief scale s varies from 0.1 (pessimistic) to 2 (optimistic) under the counterfactual scenario where prescribers inform patients about which pharmacies carry buprenorphine. The vertical line at s=1 marks baseline beliefs.

HETEROGENEITY IN PREFERENCES AND SEARCH COSTS

Figures F1a and F1b plot the distributions of the estimated deterministic preference component $(\hat{\delta})$ and the effective search cost location parameter $(\hat{\mu})$ at the patient–pharmacy level. Patients whose default pharmacy does not carry buprenorphine differ systematically from those whose default pharmacy does. The "default not available" group tends to have both lower $\hat{\delta}$ and higher $\hat{\mu}$, suggesting that simply making the default pharmacy available would not be enough to close the initiation gap for patients with intrinsically low willingness to start treatment.

These differences may reflect underlying selection. Patients with low $\hat{\delta}$ are disproportionately socially disadvantaged (Table 1) and more likely to live in neighborhoods where pharmacies face heightened risks of diversion and theft.³¹ The implication is that policies aimed solely at expanding availability are unlikely to be sufficient. A more effective approach must combine supply-side measures that reduce search costs, such as lowering uncertainty about which pharmacies carry buprenorphine, with demand-side interventions that increase patients' willingness to initiate treatment, such as modest monetary incentives. Only through this dual strategy can policy meaningfully raise initiation shares.

 $^{^{31}}$ It is also possible that pharmacies carry buprenorphine in response to local conditions: when expected demand is low, they rationally choose not to carry it.




Figure F1. Distribution of Estimated Preferences and Effective Search Costs

Note: This figure shows the distribution of estimated preferences $(\hat{\delta}_{ij})$ and effective search costs $(\hat{\mu}_{ij})$ at the patient-pharmacy level. Estimates are constructed from the coefficients in Table 1. For each outcome, the distribution is shown separately for patients whose default pharmacy stocks buprenorphine and those whose default does not.

COUNTERFACTUAL ROBUSTNESS




Figure G1. Buprenorphine Purchase Counterfactuals Under Correct Beliefs

Note: The figure is produced similar to Figure 5, but under the assumption that patients have correct beliefs about which pharmacies stock buprenorphine.




Figure G2. E-Prescription Counterfactuals Under Correct Beliefs

Note: The figure is produced similar to Figure 7, but under the assumption that patients have correct beliefs about which pharmacies stock buprenorphine.




Figure G3. Subsidy Counterfactuals Under Correct Beliefs

Note: The figure is produced similar to Figure 8, but under the assumption that patients have correct beliefs about which pharmacies stock buprenorphine.

*

References

- **Abaluck, Jason, and Abi Adams-Prassl.** 2021. "What Do Consumers Consider Before They Choose? Identification from Asymmetric Demand Responses." *The Quarterly Journal of Economics*, 136(3): 1611–1663.
- Abouk, Rahi, Rosalie Liccardo Pacula, and David Powell. 2019. "Association Between State Laws Facilitating Pharmacy Distribution of Naloxone and Risk of Fatal Overdose." *JAMA Internal Medicine*, 179(6): 805.
- Alpert, Abby, David Powell, and Rosalie Liccardo Pacula. 2018. "Supply-Side Drug Policy in the Presence of Substitutes: Evidence from the Introduction of Abuse-Deterrent Opioids." American Economic Journal: Economic Policy, 10(4): 1–35.
- **Ambokar, Sumedh, and Kian Samaee.** 2019. "Inaction, search costs, and market power in the us mortgage market." *University of Pennsylvania mimeo*.
- **Armstrong, Mark.** 2017. "Ordered Consumer Search." *Journal of the European Economic Association*, 15(5): 989–1024.
- Choi, Michael, Anovia Yifan Dai, and Kyungmin Kim. 2018. "Consumer Search and Price Competition." *Econometrica*, 86(4): 1257–1281.
- Compiani, Giovanni, Gregory Lewis, Sida Peng, and Peichun Wang. 2024. "Online Search and Optimal Product Rankings: An Empirical Framework." *Marketing Science*, 43(3): 615–636.
- Dever, Jill A., Marci F. Hertz, Laura J. Dunlap, John S. Richardson, Sara Beth Wolicki, Bradley B. Biggers, Mark J. Edlund, Michele K. Bohm, Didier Turcios, Xinyi Jiang, Hong Zhou, Mary E. Evans, and Gery P. Guy. 2024. "The Medications for Opioid Use Disorder Study: Methods and Initial Outcomes From an 18-Month Study of Patients in Treatment for Opioid Use Disorder." Public Health Reports®, 139(4): 484–493.

- Donahoe, J. Travis. 2025. "Supplier Enforcement and the Opioid Crisis." Working Paper.
- Evans, William, Ethan Lieber, and Patrick Power. 2018. "How the Reformulation of OxyContin Ignited the Heroin Epidemic." National Bureau of Economic Research w24475, Cambridge, MA. Working Paper.
- **Gui, Lance.** 2024. "Closing the Opioid Use Disorder Treatment Gap: Expanding Nurse Practitioners' Prescriptive Authority." SSRN Electronic Journal. Working Paper.
- **Handel, Benjamin R.** 2013. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review*, 103(7): 2643–2682.
- Hastings, Justine, Thomas Kane, and Douglas Staiger. 2006. "Preferences and Heterogeneous Treatment Effects in a Public School Choice Lottery." National Bureau of Economic Research w12145, Cambridge, MA. Working Paper.
- Heiss, Florian, Daniel McFadden, Joachim Winter, Amelie Wuppermann, and BoZhou. 2021. "Inattention and Switching Costs as Sources of Inertia in Medicare Part D."American Economic Review, 111(9): 2737–2781.
- Ho, Kate, Joseph Hogan, and Fiona Scott Morton. 2017. "The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program." The RAND Journal of Economics, 48(4): 877–905.
- **Hong, Han, and Matthew Shum.** 2006. "Using Price Distributions to Estimate Search Costs." *The RAND Journal of Economics*, 37(2): 257–275.
- **Honka, Elisabeth.** 2014. "Quantifying Search and Switching Costs in the US Auto Insurance Industry." *The RAND Journal of Economics*, 45(4): 847–884.
- Hortaçsu, Ali, Seyed Ali Madanizadeh, and Steven L. Puller. 2017. "Power to Choose? An Analysis of Consumer Inertia in the Residential Electricity Market." *American Economic Journal: Economic Policy*, 9(4): 192–226.
- Kim, Jun B., Paulo Albuquerque, and Bart J. Bronnenberg. 2010. "Online Demand Under Limited Consumer Search." *Marketing Science*, 29(6): 1001–1023.

- Kim, Jun B., Paulo Albuquerque, and Bart J. Bronnenberg. 2017. "The Probit Choice Model Under Sequential Search with an Application to Online Retailing." *Management Science*, 63(11): 3911–3929.
- Krawczyk, Noa, Bianca D. Rivera, Victoria Jent, Katherine M. Keyes, Christopher M. Jones, and Magdalena Cerdá. 2022. "Has the Treatment Gap for Opioid Use Disorder Narrowed in the U.S.?: A Yearly Assessment from 2010 to 2019"." International Journal of Drug Policy, 110: 103786.
- Larochelle, Marc R., Dana Bernson, Thomas Land, Thomas J. Stopka, Na Wang, Ziming Xuan, Sarah M. Bagley, Jane M. Liebschutz, and Alexander Y. Walley. 2018. "Medication for Opioid Use Disorder After Nonfatal Opioid Overdose and Association With Mortality: A Cohort Study." Annals of Internal Medicine, 169(3): 137–145.
- Luco, Fernando, Benjamin Klopack, and Eric Lewis. 2024. "Rebuilding After the Storm: Firm Turnover and Consumer Welfare After Hurricane Harvey." SSRN Electronic Journal. Working Paper.
- Ma, Jun, Yan-Ping Bao, Ru-Jia Wang, Meng-Fan Su, Mo-Xuan Liu, Jin-Qiao Li,
 Louisa Degenhardt, Michael Farrell, Frederic C. Blow, Mark Ilgen, Jie Shi, and
 Lin Lu. 2019. "Effects of Medication-Assisted Treatment on Mortality among Opioids Users:
 A Systematic Review and Meta-Analysis." Molecular Psychiatry, 24(12): 1868–1883.
- Marcum, Zachary A., Julia Driessen, Carolyn T. Thorpe, Walid F. Gellad, and Julie M. Donohue. 2014. "Effect of Multiple Pharmacy Use on Medication Adherence and Drug-Drug Interactions in Older Adults with Medicare Part D." Journal of the American Geriatrics Society, 62(2): 244–252.
- Moraga-González, José Luis, Zsolt Sándor, and Matthijs R Wildenbeest. 2023. "Consumer Search and Prices in the Automobile Market." The Review of Economic Studies, 90(3): 1394–1440.
- Rees, Daniel, Joseph Sabia, Laura Argys, Joshua Latshaw, and Dhaval Dave. 2017. "With a Little Help from My Friends: The Effects of Naloxone Access and Good Samaritan

- Laws on Opioid-Related Deaths." National Bureau of Economic Research w23171, Cambridge, MA. Working Paper.
- Santos, Babur De Los, Ali Hortaçsu, and Matthijs R Wildenbeest. 2012. "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior." American Economic Review, 102(6): 2955–2980.
- Severtson, Stevan Geoffrey, Becki Bucher Bartelson, Jonathan M. Davis, Alvaro Muñoz, Michael F. Schneider, Howard Chilcoat, Paul M. Coplan, Hilary Surratt, and Richard C. Dart. 2013. "Reduced Abuse, Therapeutic Errors, and Diversion Following Reformulation of Extended-Release Oxycodone in 2010." The Journal of Pain, 14(10): 1122–1130.
- Skains, Rachel M., Lindy Reynolds, Nicholas Carlisle, Sonya Heath, Whitney Covington, Kyle Hornbuckle, and Lauren Walter. 2023. "Impact of Emergency Department-Initiated Buprenorphine on Repeat Emergency Department Utilization." Western Journal of Emergency Medicine, 24(6).
- **Soliman, Adam.** 2023. "Disrupting Drug Markets: The Effects of Crackdowns on Rogue Opioid Suppliers." SSRN Electronic Journal.
- Sullivan, Ross W., Laura M. Szczesniak, and Susan M. Wojcik. 2021. "Bridge Clinic Buprenorphine Program Decreases Emergency Department Visits." *Journal of Substance Abuse Treatment*, 130: 108410.
- Ursu, Raluca M. 2018. "The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions." *Marketing Science*, 37(4): 530–552.
- Weiner, Scott G., Dima M. Qato, Jeremy Samuel Faust, and Brian Clear. 2023. "Pharmacy Availability of Buprenorphine for Opioid Use Disorder Treatment in the US." *JAMA Network Open*, 6(5): e2316089.
- Weitzman, Martin L. 1979. "Optimal Search for the Best Alternative." *Econometrica*, 47(3): 641.

Yarborough, Bobbi Jo H., Scott P. Stumbo, Shannon L. Janoff, Erin M. Keast, Michael C. Leo, and Sarah J. Leitz. 2024. "Reduced Emergency Department Use among Insured Individuals Receiving Extended-Release Buprenorphine in a Health System Setting." Drug and Alcohol Dependence Reports, 11: 100233.