Do Pharmacies Matter?

Joseph Battles*
University of Chicago

October 31, 2025
Click Here for Latest Version

Abstract: Despite potential health benefits of prescription drugs, many patients do not take them as prescribed. While prior work has emphasized the role of prices and insurers in patients' drug consumption decisions, this paper provides novel evidence on a non-price driver of drug consumption: pharmacies. Using a staggered event study design, I find that pharmacy closures cause Medicare patients' drug use to decline in the short-run—especially among low-income patients—but increase in the long run. To explain the long-run increase in drug consumption following a pharmacy closure, I model three potential mechanisms driving the reduced-form effect: temporary disruption/switching frictions, permanent changes in patient costs (e.g., copays or travel distance), and permanent shifts to higher- or lower-dispensing pharmacies. To quantify the relative impacts of these mechanisms, I estimate a two-way-fixed-effects model in the style of Abowd, Kramarz, and Margolis (1999) of pharmacies' effects on low-income patients' drug use. Combining the pharmacy effects from the AKM model with my reduced-form closure analyses, I find that the long-run increase in drug consumption following a pharmacy closure is explained by patients switching from lower-dispensing pharmacies (which are disproportionately likely to close) to higher-dispensing pharmacies. More generally, the variation in pharmacy fixed effects is about half that of prescriber fixed effects, indicating that pharmacies matter for drug consumption.

*jbattles@uchicago.edu I am immensely indebted to Manasi Deshpande, Joshua Gottlieb, Matthew Notowidigdo, and Bradley Shapiro for their advising and guidance throughout the life of this project. Zarek Brot, Sara Drango, Deniz Dutz, Thibaut Lamadon, Camilla Schneier, Maggie Shi, and Haruka Uchida also provided extensive and helpful feedback. Kristin Hitchcock and Mohan Ramanujan provided invaluable assistance with accessing and using Medicare claims data. Many thanks to the participants in UChicago/Northwestern T32 HSR Workshop, AEA CSQIEP Workshop, UChicago Third-Year Ph.D. Seminar, Health Care Reading Group, Public/Labor Lunch, and Student Applied Micro Lunch for useful discussions and feedback. I gratefully acknowledge support from the Becker Friedman Institute and the Agency for Healthcare Research and Quality under grant award T32 HS000084 (PI: Tamara Konetzka, Ph.D.). This paper's contents are solely the responsibility of the author and do not necessarily represent the official views of the AHRQ.

1 Introduction

Prescription drugs can have large positive impacts on health, and many have been shown to be cost-effective. Buxbaum et al. (2020) argue that prescription drugs drove 35% of the decline in U.S. mortality between 1990 and 2015. Drugs can also prevent more costly types of care, such as hospitalizations, which for publicly insured patients, are largely paid by taxpayers through Medicare and Medicaid. Indeed, (Heller et al., 2017) estimate that universal use of certain cheap, high-value medications by all adults of a certain age would save money.

One barrier to realizing these benefits is patients not taking drugs as directed (Osterberg and Blaschke, 2005; Cutler et al., 2018). For this reason, increasing high-value drug use has been a policy priority. These policies include reducing or eliminating copays for preventive and high-value drugs (e.g., vaccines, insulin), and imposing reimbursement floors on pharmacy payments to prevent pharmacy exits.[1] Research on increasing prescription drug consumption has focused on drug prices faced by patients (copays), and more recently, on targeted, insurer-levied hassles (e.g., prior authorization requirements). However, far less is known about the non-price drivers of patient drug consumption decisions, especially the "last-mile" problem of prescribing and physically dispensing prescription drugs to patients.

In this paper, I provide some of the first evidence on a non-price driver of drug consumption: pharmacies. Pharmacies are among the most commonly used medical providers in the U.S. health care system; on average, patients visit pharmacies twice as often as a doctor Berenbrok et al. (2020); Valliant et al. (2022). However, there is little evidence about whether they matter for patient behavior. I start by estimating the reduced-form effect of pharmacy closures on drug dispensing. I then combine these estimates with an AKM-style model of patients' prescription drug consumption. The combination of these analyses allows me to separate the mechanisms through which pharmacy closures affect patient drug consumption: temporary switching costs, changes in drug costs (e.g., travel distance or copays), and reallocation to pharmacies with different time-invariant drug dispensing propensities. I find that pharmacy closures cause temporary declines in patients' prescription drug use, but long-term *increases* within a year of the closure. The AKM

[1] For example, the Inflation Reduction Act eliminated vaccine copays and capped insulin copays for Medicare recipients. States that have imposed reimbursement floors include Kentucky, Tennessee, West Virginia, and Alabama.

model reveals substantial heterogeneity in pharmacies' effects on drug consumption. Combining the model-estimated pharmacy effects with the reduced-form closure analyses, I show that higher rates of exit among low-dispensing pharmacies—resulting in patients switching to higher-dispensing pharmacies—account for the long-run increases in drug use.

I start by building a conceptual model of patients' pharmacy choice and prescription drug consumption. This model captures three key factors determining how a pharmacy closure impacts drug use: changes in costs, reallocation effects from changing to a higher or lower-dispensing pharmacy, and disruption/switching frictions. These three mechanisms can have opposing and time-varying impacts on patients' medication use: switching costs are temporary and reduce drug consumption, while changes in patient costs and reallocation effects are permanent and can either increase or decrease consumption. I show that reduced-form estimates of the effect of pharmacy closures recovers the combined effect of all three mechanisms. A two-way-fixed-effects model in the style of AKM can, under certain assumptions, recover pharmacy dispensing effects, which I can then use to estimate the impacts of the reallocation mechanism.

In the reduced-form analysis, I first document high rates of pharmacy closure: 2% close each year (Guadamuz et al., 2020). Because patients visit pharmacies far more often than other health care providers (Berenbrok et al., 2020), pharmacy exits impact many more patients than other types of provider exits. Among Medicare patients specifically, I show that over 20% will experience at least one pharmacy exit during their tenure on Medicare.

Using Medicare claims data, I estimate the causal effects of pharmacy closure using a difference-in-difference (DiD) design. I implement a stacked DiD using pharmacy closures one year later as controls for pharmacies closing today. The stacked design addresses recent concerns about identification in designs with staggered treatment and relies on the standard parallel trends assumption (Deshpande and Li, 2019; Baker, Larcker and Wang, 2022). In addition to flat pre-trends, I present robustness checks that reinforce the validity of this assumption.

The DiD regressions show immediate reductions in patients' drug consumption following pharmacy closure, followed by long-run increases. Among patients previously shopping at the closing pharmacy, prescription drug use falls by 1.1% in the quarter of closure. The declines impact all drug use, including medications typically categorized as high-value,[2] and disproportionately impact low-income patients. Further heterogeneity analyses reveal that exits thought to make access to drugs especially inconvenient to acquire, such as those in rural areas or that create “pharmacy deserts”, are no worse than the average closure. Most notably, the sign of the treatment effect reverses within a year and drug consumption significantly increases as a result of pharmacy closure.

What explains the long-run increase in drug consumption following a pharmacy closure? Based on my conceptual model, I hypothesize that reallocations to higher-dispensing pharmacies drive the long-run increases. Testing this hypothesis requires an estimate of the effect of each pharmacy on patients' drug consumption. I obtain these estimates using a two-way fixed effects model in the style of Abowd, Kramarz and Margolis (1999). The model is identified using low-income patients who move pharmacies,[3] The AKM model shows that pharmacies impact patients' drug consumption, explaining about 1% of the variance in patients' prescription drug use. To contextualize this number, I estimate a separate model of the role of prescribers, which prior work has argued have substantial effects on patient drug use (Laird and Nielsen, 2016), as a benchmark. I find that pharmacies' impact is about half as large as that of prescribers, indicating that pharmacies matter for patient behavior.

To test the reallocation hypothesis, I take these pharmacy fixed effects back to the reduced-form closure analyses and find that low-income patients disproportionately move to higher-dispensing pharmacies post-closure. Since the pharmacy fixed effects are estimated from a separate set of patients as the closure analysis, this result does not reflect any mechanical relationship. Moves from low- to high-dispensing pharmacies explain the long-run increases in low-income patients' drug use. Connecting these pieces back to the conceptual model, I show that disruption costs depress drug use in the short-run, but the reallocation effect dominates in the long-run and explains the long-run increase in drug consumption.

Using my reduced-form and AKM estimates, I evaluate a 2022 policy change in California's Medicaid program that increased payments for small pharmacies. This policy is one example of

[2]For instance, using the classification developed by Chandra, Gruber and McKnight (2010), discussed more later.

[3]Prior work has shown that prices vary at the plan-pharmacy-year level and cause sorting (Starc and Swanson, 2021). This sorting causes match effects that violate assumptions required to estimate AKM. Low-income Medicare patients face much less variation in drug prices, thus addressing this issue.

many recent policies that may that could prevent pharmacy closures and thereby improve access to prescription drugs. However, if this policy prevented the average closure, it would *reduce* low-income patients' drug use by 1.1% in the long-run, despite back-of-the-envelope calculations suggesting the policy cost millions of dollars per closure averted. This finding highlights the importance of heterogeneity in dispensing efficacy across pharmacies, and the effect of this heterogeneity on the effectiveness of policies to improve access to prescription drugs.

The paper proceeds as follows. Section 1.1 situates the main findings within the literature. Section 2 provides the institutional details and describes the data, while Section 3 develops a conceptual model of patients' drug consumption. Section 4 presents the main reduced-form analyses of the effects of pharmacy exits. Section 5 describes and estimates the AKM model and incorporates these back into the reduced form. Section 6 concludes.

1.1 Literature and Contributions

Unlike prior work focused on prices or targeted insurer-imposed hassles, I focus on non-price costs arising from the "last-mile problem" of drug prescribing and dispensing at pharmacies.

Although higher prices reduce drug consumption, low price elasticities and patients' tendency to skip valuable drugs highlight the need to consider non-price costs. Studies have consistently found strictly negative price elasticities for prescription drugs across broad patient groups, including the privately-insured (Chandra, Gruber and McKnight, 2010), high-income Medicare recipients (Starc and Swanson, 2021; Chandra, Flack and Obermeyer, 2024) and low-income Medicare patients also enrolled in Medicaid (Hollrah, 2024). However, demand is inelastic and patients often do not take high-value drugs, even when free (Einav, Finkelstein and Polyakova, 2018; Choudhry et al., 2011).

Recent evidence has shown large effects of insurer-levied non-price costs, but these barriers are designed to reduce consumption of either narrow sets of patients or expensive/low-value drugs. Layton et al. (2022), Brot-Goldberg et al. (2023), Brot et al. (2024), and Agafiev Macambira et al. (2025) show that quantity limits, formulary restrictions, and prior authorization requirements all cause patients to drastically reduce drug consumption.[4] However, these policies cannot address inefficiently low consumption of high-value drugs (Choudhry et al., 2011; Chandra, Flack and

[4] For example, Brot et al. (2024) shows that prior authorization requirements reduce use of targeted drugs by 25%.

Obermeyer, 2024).[5] Reducing non-price costs of valuable medications may be especially important in cases where consumption is inefficiently low even at prices close to zero (i.e., low-income Medicare and Medicaid recipients' copays).[6] In contrast, pharmacy closures impact all drugs and a wide range of patients, and the broader non-price mechanisms I study can *increase* drug consumption.

My empirical framework and rich Medicare data allow addressing data and identification challenges that limit causal interpretations of a small literature examining pharmacy exits. Qato et al. (2019) and Anderson et al. (2024) find large, permanent negative associations between pharmacy exits and patients' drug use. However, the former suffers from a substantial non-random missing data issue, while the latter examines a narrow set of drugs and a small number of pharmacy exits in a single state (Colorado).[7]

The richness of the Medicare data also allows me to estimate two-way fixed effect models to estimate pharmacies' broader role in patients' drug consumption. This work builds on a large literature in labor using these models to estimate firm effects on wages (Abowd, Kramarz and Margolis, 1999; Card, Heining and Kline, 2013; Bonhomme, Lamadon and Manresa, 2019), and a more recent set of papers applying these models to health care settings (Finkelstein, Gentzkow and Williams, 2016; Badinski et al., 2023; Mourot, 2025). I also build on this prior work by using the model estimates inside my reduced-form analyses, allowing me to disentangle underlying mechanisms that cannot be identified using the reduced-form alone.

While there is limited evidence on pharmacy exits specifically,[8] evidence from other health care settings has found that effects of changes in the supply of providers depend on the types of providers entering or exiting, the degree of switching costs, and the substitutes available to patients. For instance, Fischer, Royer and White (2023) and Olenski (2022) find *positive* effects of provider exits on patient health due to substitution to higher-quality providers, while Carroll (2023), Sabety (Forthcoming) and Garthwaite, Gross and Notowidigdo (2018) find the opposite. Thus, even the sign of the effect of pharmacy exit on patients' drug consumption is an open empirical question.

[5] This inefficiency can arise either from patients' behavioral biases or fiscal externalities from increases in costly, insured types of care caused by drug non-compliance that insuers cannot contract over.

[6] Cheap drugs that reduce costly heart attacks and vaccines are two examples of drugs where positive externalities motivate efforts to lower non-price costs even when prices are zero (Choudhry et al., 2011; Chevalier et al., 2022).

[7] One additional unpublished paper, Gannaway (2019), finds a null effect of exits, but examines only 17 and admits that many are likely false positives not reflecting true exits.

[8] Janssen and Zhang (2023) examines pharmacies' impacts on patient drug consumption, but only of opioids.

2 Institutional Details and Data

2.1 What do pharmacies do?

Obtaining a prescription medication in the United States requires many steps. First, the patient needs an order for a specific drug, dosage, and quantity from a licensed physician (the "prescriber").[9] Second, this order is transferred to a pharmacy (the "dispenser"), where a pharmacist prepares the medication. The patient must then travel to the pharmacy or arrange to have the drug mailed or delivered. The pharmacy then imposes administrative hurdles (e.g., prior authorization) and collects any copayment required by the patients' insurer before actually providing the medication to the patient. Finally, the patient then decides whether to consume the drug and if so, whether to refill. At this point, the patient may just return to the pharmacy if the prescription has refills remaining, but otherwise must obtain a new prescription from the prescriber. Importantly, the prescriber and dispenser are usually different entities located in different places.

While typically lacking the authority to prescribe medication or set prices for insured patients,[10] pharmacies might still have a major impact on whether a patient takes a drug as directed. As the actual dispenser of medications, pharmacies interact frequently with patients and often must coordinate with a patient's doctors and insurer. For instance, they are responsible for collecting insurer authorizations when required, and monitoring for harmful drug interactions (especially when a patient gets prescriptions from multiple doctors). In this role, pharmacists can help patients streamline their prescription fills, identify cheaper or more convenient options (e.g., generics or mail-order/delivery options), and help them navigate interactions with prescribers and insurers to change or refill their prescriptions. More broadly, the pharmacy retail environment can also impact patients' behavior. These factors range from stock outs, long lines, and the helpfulness of the staff, to how well the pharmacy manages and reminds patients to refill and pickup medications. Finally, pharmacists are trained to monitor for drug abuse, harmful dosages, and dangerous drug interactions, as well as counsel patients on correct administration of the drugs (e.g., timing or whether the drug should be taken with food).

[9] Depending on the state and the drug, certain non-physician health professionals may have prescribing authority, such as physician assistants, nurse practitioners, or dentists. However, for the sake of simplicity, I use the term "physician" or "doctor" to refer to all of these prescribers.

[10] Some states allow pharmacists to prescribe a very limited set of drugs directly (e.g., certain vaccines).

Drug dispensing accounts for a significant share of drug spending, especially for generic drugs. Many widely prescribed and highly-effective generics have wholesale prices below a dollar for a 30-day supply, but cost pharmacies many times more to dispense.[11]

2.2 Medicare Drug Benefits

Medicare patients account for a large share of U.S. prescription drug consumption. They receive drug benefits through an optional program ("Part D"), which began in 2006. Because Medicare eligibility is restricted to elderly and certain disabled adults, Part D patients are disproportionately sick, accounting for 14% of the population, but 31% ($105B) of national drug expenditures in 2019.[12]

Medicare Part D is provided through private insurers in a heavily-regulated marketplace of plans. Insurers compete to enroll patients and must meet certain standards set by Medicare over drugs covered and the amount of cost-sharing required of patients. Importantly, while Medicare requires Part D insurers provide "adequate" access to pharmacies, the government otherwise imposes few constraints on contracting between insurers, pharmacies, and manufacturers. The lack of constraints on pharmacy payments in Part D is similar to many other commercial insurance settings and is one motivation for recent policy efforts to increase pharmacy reimbursements.

One key way the government influences patient and insurer behavior in this context is through large subsidies, especially for low-income patients. Overall, the federal government pays about 75% of the costs of Part D, with patients covering the remainder through a combination of premiums and copayments. Copayments for a given drug on a given plan can vary depending on the pharmacy, with many plans offering lower prices at a preferred subset of pharmacies (Starc and Swanson, 2021). Extra support for low-income patients is typically provided through the "low-income subsidy" (LIS) which replaces Medicaid drug coverage and is available to Medicare enrollees earning below 150% of the federal poverty line (28% of Part D recipients in 2019).[13] LIS patients are automatically enrolled in a plan if they do not select one themselves and typically pay no premiums and low or no copayments for covered drugs. Importantly, these copayments are administratively set by congress

[11] For example, pharmacies' national average drug acquistion cost (NADAC) for the three most commonly prescribed drugs in Medicare Part D in 2022 ranged from $0.30 to $3 for a 30-day supply, while pharmacy dispensing fees in most state Medicaid programs range from $5-15.

[12] Enrollment number from KFF and spending number from CMS National Health Expenditure Data.

[13] See the following KFF article for details.

and are typically zero or very low. For example, Congress set copays for Medicare patients earning less than the Federal Poverty Line to $1 per generic drug fill in 2006.

2.3 Medicare Data

The backbone of my empirical analyses is the Medicare Part D Event (PDE) File. This individual-level administrative dataset includes a random 20% longitudinal sample of Part D patients from 2006-2019. The PDE file includes all drug claims for the 20% sample throughout their tenure in Medicare.[14] Each row describes one drug claim, including the fill date, the exact drug and quantity dispensed, and identifiers for the patient, dispenser, prescriber, and insurer.

Using patient identifiers in the PDE file, I merge the drug data with other Medicare files to examine patient characteristics and outcomes beyond prescription drug use. The Master Beneficiary Summary File (MBSF) includes key patient demographics, including ZIP code of residence, LIS enrollment, date of birth/death, race, and gender. The Medicare Provider Analysis and Review File (MedPAR) tracks patients' hospitalization and nursing home stays. It provides detailed usage and spending data in these settings, which allows me to study spillover effects of pharmacy closures on other types of health care use.

Finally, I use other data sources to extend certain analyses beyond the last year of my claims data (2019). These include public data from Medicare about all Part D pharmacies and which insurers they contract with in each year, along with public plan-by-year-level data on plan characteristics and patient enrollment levels. Importantly, these datasets extend through 2025, which I use to study pharmacy entries and exits that I am unable to examine with the claims data.

2.4 External Data Linkages to Define Pharmacies

Unlike patient identifiers, pharmacy and prescriber identifiers can be linked to external data, which I use to define pharmacy closures and the patients affected by them.

Each Part D claim includes an identifier for the dispenser of the prescription drug. From 2009 onward, 99% of claims include the national provider identifier (NPI) of the dispenser.[15] NPIs can

[14] I drop 2006 from all analyses because Part D was rolled out during this year. In 2019, the PDE file included 301 million drug claims from 9 million patients.

[15] For 2007 to 2009, many claims used a different identifier from the National Council of Prescription Drug Providers,

change over time and a given location can have multiple NPIs at once,[16] making it challenging to use NPIs alone to define pharmacy entry or exits. I thus link the NPIs provided in the claims data to the public National Plan and Provider Enumeration System (NPPES), which provides a physical address for each NPI and classifies the NPI as an individual or institution.

To estimate the effects of real exits of physical pharmacy locations, I make three sample restrictions. First, I subset to NPIs classified as an institution (rather than individual) and group NPIs listed at the same address into a single identifier, which I refer to as a "harmonized NPI."[17] I define entry and exit as the first and last month that a harmonized NPI fills a Part D claim. This definition of exit excludes ownership changes or pharmacies moving location without changing NPIs. Second, to avoid conflating effects of entry and exit, I restrict the closure sample to include only pharmacies that enter at least 24 months prior to exit. Finally, I use a dispensing threshold to exclude a small number of exits that are likely mail-order pharmacies. By definition, patients do not travel to these pharmacies, so the meaning of a location "exiting" is ill-defined.[18] Note that the restrictions above are used only to define the set of closing pharmacies. All outcome measures are based on the full sample of Part D data. For instance, while mail-order pharmacies are excluded from the exiting sample, if patients are more likely to switch to mail order after their physical pharmacy closes, this substitution will be captured in the outcomes. I provide more details about this process in Appendix B.

3 Conceptual Model

3.1 Choice of Prescription Drug Consumption Conditional on Pharmacy

I develop a simple model of patients' drug consumption decisions that highlights the potential role played by pharmacies. Consider a patient i choosing a quantity of prescription drugs to consume y to maximize utility in time period t. I start by fixing the patient's pharmacy, $j(i,t)$; conditional on this choice, a patient's utility $u_{ijt}(y)$ is a function of a pharmacy effect (q_j), patient effect (α_i),

but from 2009 onwards, 99% of claims listed an NPI. More details about addressing other IDs are in Appendix B.

[16]Importantly, however, the reverse is not true. Even if under common ownership, CMS typically requires that each location of a pharmacy have its own NPI.

[17]In rare cases, doctors dispense drugs directly, in which case the dispensing NPI could be an individual. I do not consider these dispensers as pharmacies, but count these fills when measuring patient prescription drug use.

[18]For example, if a large mail-order pharmacy closed one distribution center and just fulfilled the orders from a different warehouse, patients are unlikely to even be aware of the "exit".

health shock (h_{it}), and costs ($C_{ijt}(y)$). (I suppress the arguments of j for notational simiplicity.) The pharmacy effect captures time- and patient-invariant pharmacy characteristics that increase (e.g., refill reminders) or decrease (e.g., longer lines) patients' drug consumption. Adapting Badinski et al. (2023), the health shock represents the time-varying optimal drug consumption level based on some objective threshold, while the patient fixed effect reflects misperceptions or preferences that differ from this level.[19] These two components comprise the benefits a patient receives from drug consumption, with deviations from the optimal level parametrized as imposing a quadratic reduction in utility (again, following Badinski et al. (2023)). Costs include price and non-price costs, such as travel distance. I parameterize utility as shown in Equation 1.

$$u_{ijt}(y) = \underbrace{\alpha_i y - \frac{1}{2}(y - h_{it})^2}_{\text{benefits}} - \underbrace{C_{ijt}(y)}_{\text{costs}} + \underbrace{q_j y}_{\text{pharm. effect}} \tag{1}$$

Under the assumption that costs increase in y, Equation 2 shows the implicit function from the first-order condition that determines utility-maximizing drug consumption, y^*_{ijt}. If costs are linear ($C_{ijt}(y) = c_{ijt}y$), then the expression simplifies to Equation 3.

$$y^*_{ijt} = \alpha_i + h_{it} - C'_{ijt}(y^*_{ijt}) + q_j \tag{2}$$

$$y^*_{ijt} = \alpha_i + h_{it} - c_{ijt} + q_j \tag{3}$$

Consider a case where patient i switches from pharmacy j to j' in period 2. The observed change in drug consumption can be decomposed into the three parts shown in Equation 4. Note that the time-invariant patient fixed effect drops out, leaving the change in the health shock, costs, and the pharmacy fixed effects. While drug consumption is observed, the terms on the right-hand side may not be. For instance, if pharmacy changes are correlated with unobserved health shocks, then my estimates of the pharmacy fixed effects will be biased.

$$y^*_{i2} - y^*_{i1} = (h_{i2} - h_{i1}) - (c_{ij'2} - c_{ij1}) + (q'_j - q_j) \tag{4}$$

[19] For instance, Chandra, Flack and Obermeyer (2024) find that patients underestimate health consequences of drug discontinuation.

3.2 Choice of Pharmacy

Patients choose a pharmacy to maximize utility from drug consumption and non-drug pharmacy characteristics. I capture these non-drug characteristics with a non-drug pharmacy fixed effect (ζ_j), which does not scale with drug consumption. Patients receive a type 1 extreme value preference shock in each period (γ_{ijt}), and choose a pharmacy $j^*_{it} \in J(t)$ that satisfies Equation 5, where y^*_{ijt} is the solution to the patient's drug utility maximization problem (Equation 1). The shock captures any changes in patients' preferences over non-drug pharmacy characteristics, such as the pharmacy brand or the products sold at the pharmacy other than prescription drugs. Patients' outside option is to consume no drugs. In this case, consumption and utility both equal 0.

$$j^*_{it} = \underset{j \in J(t)}{\arg\max} \left[u_{ijt}(y^*_{ijt}) + \zeta_j + \gamma_{ijt} \right] \tag{5}$$

Under the assumption of linear costs, there are two reasons a patient changes their pharmacy: a change in costs or the preference shock.[20] A change in costs could reflect prices changing, but also patients moving homes, which would change their pharmacy travel costs. If these shocks are independent of the health shock, then patients changing pharmacies can be used to identify the effects of costs and pharmacies on patients' drug use.

3.3 Adapting Model to Study Pharmacy Closures

I now adapt the model to examine the effects of a patient's pharmacy closing. The parameter of interest is shown in Equation 6.

$$\mathbb{E}\left[y^*_{i,t} - y^*_{i,t=-1} \mid j(i, t=-1) = \tilde{j}, \tilde{j} \notin J(t \geq 0) \right] \tag{6}$$

I then substitute in Equation 4, with one modification. Because closures typically occur abruptly and force patients to change pharmacies, I add a time-varying disruption cost s_t (with $s_{t<0} = 0$). (Note that s_t can just be thought of as an additional component of $C_{ijt}(y)$ defined in Equation 1.)

[20] Without the functional form assumption on costs, the health shock could also cause patients to switch pharmacies if the relative costs change at higher or lower levels of drug consumption (i.e., if costs are non-linear).

Equation 7 shows the new expression after the substitution and modification.

$$\mathbb{E}\left[\underbrace{(h_{it} - h_{i,t=-1})}_{\text{health shock}} - \underbrace{(c_{i,j',t} - c_{i,j,t=-1}) + (q'_j - q_j) - s_t}_{\text{closure effect}} \mid j(i, t = -1) = \tilde{j}, \tilde{j} \notin J(t \geq 0)\right] \tag{7}$$

The terms inside the "closure effect" underbrace comprise the reduced-form parameter of interest. Thus, the key challenge is to separate this effect from the health shock. As described in the next section, the stacked difference-in-difference design provides a plausible way to control for the health shock and recover the closure effect among the set of patients experiencing a pharmacy closure. Importantly, note that the effect of closures on drug consumption is ambiguous in the model.[21] In particular, the sign of the effect of closures depends on the direction and relative magnitude of the three constituent components of the closure effect. The reduced-form DiD will only capture the overall average of the closure effect, while the two-way fixed effects model presented later will allow me to separate out the pharmacy fixed effects.

4 Pharmacy Closures

4.1 Descriptive Facts about Pharmacies and Pharmacy Closures

Before presenting the reduced form regressions and results, I first provide some descriptive facts illustrating that pharmacy entry and exits are common. Panel (a) of Figure 1 shows the annual number of pharmacy entries and exits between 2009 and 2019.[22] Pharmacy turnover is common: there are over 2,000 entries and exits per year on average and 23% of pharmacies open in 2009 had closed by 2019.[23] While the total number of pharmacies increased due to higher rates of entry through 2017, closures grew more common throughout the study period, with the net number of pharmacies beginning to decline in 2018. These closures impacted over 3.5 million Part D patients

[21] The effect on welfare is weakly negative by revealed preference, but the fact that patients have terms unrelated to drug use in their utility function (Equation 5) means that drug use may increase.

[22] Pharmacies are defined using the criteria described above in Section 2.4, except that I relax the requirement that a pharmacy be open for at least two years. (Otherwise, there would mechanically be no entries in 2018 or 2019.) Because there were pharmacies switching from one identifier to another in 2007 and 2008, I drop these years to avoid overstating the number of entries. More details in Appendix B.

[23] Restricting to the set of pharmacies open for at least two years, I estimate that 13,800 of the 59,900 pharmacies open in 2009 had closed by 2019.

during my sample period, including over 15,000 every month, as shown in Panel (b) of Figure 1.[24] My claims data end in 2019, but using public Medicare data, Appendix Figure A1 shows that closures have grown even more common, with a net decline of 10% of pharmacies between 2018 and 2024. Appendix Figures A2 and A3 illustrate that exits are rarely replaced by a nearby entry and that there is wide spatial heterogeneity in rates of entry and exit.

4.2 Effects of Pharmacy Exit in the Raw Data

Figure 2 illustrates the main reduced-form results with just the raw data. This figure shows monthly prescription drug consumption (measured in total days supply of medication) for two groups of patients: those experiencing a closure in month 0 (orange diamonds) and those experiencing closure one year later in month 12 (gray circles). Exactly in month 0, the patients experiencing closure see a sharp drop in medication use, while there is no similar drop among patients experiencing a closure later. However, by 6 months after closure, drug consumption has recovered among the month 0 group and continues growing faster than the later closing group through the end of the study period. Despite these differences in the months 0 through 11, the groups' consumption is remarkably similar in both levels and trends in the preceding months. The remainder of this section develops an econometric model that captures these results more formally.

4.3 Defining Treatment and Outcomes

The population most directly impacted by a closing pharmacy is the set of patients who shopped there prior to closure, but patients shopping there just before exit may be selected. For example, if a pharmacy announces that it is going to close in a month, attentive patients may immediately move their prescription to another pharmacy, leaving behind a disproportionately inattentive group of patients shopping at the pharmacy just before closure. I thus define a patient as treated by an exit if they fill at least one Part D claim at the pharmacy 7-12 months before closure. For patients who experience multiple exits during their Medicare tenure, I use only the first.

While pharmacy closures are likely non-random, I leverage plausibly exogenous variation in the exact timing of exit to motivate using patients experiencing closure in the future as a control

[24] "Impacted" means treated using the definition I present in Section 4.3.

group. I create the control group by randomly matching each treated patient to a patient that also fills at least one Part D claim in the same six month window, but who experiences a pharmacy closure exactly 12 months later.[25] Thus, the control patients fill at least one Part D prescription 19-24 months prior to the exit they experience.

My primary outcome of interest is patients' prescription drug consumption.[26] I measure consumption by summing each patients' total monthly days supply of medications across all Part D claims. Because claims can include multiple months supplied at once, I split claims of more than a 30-day supply across the subsequent months (e.g., a 90-day fill in January is treated equivalently to three 30-day fills in January, February, and March). Like all claims data and virtually all administrative data, I observe only whether the patient receives the medication, not whether she actually ingests it. However, because acquiring medications imposes non-price costs and often a copay, fills are a reasonable and commonly used proxy for consumption.[27]

4.4 Identification, Empirical Specification, and Tests of Validity

To recover the causal effects of pharmacy closures, I implement a stacked difference-in-difference (DiD) design matching patients experiencing closure in one month to control patients experiencing a closure exactly 12 months later. I estimate the model shown below in Equation 8.

$$Y_{ikt} = \delta \texttt{treated}_{ikt} + \sum_{\tau \neq -1} D^{\tau}_{ikt} + \sum_{\tau \neq -1} \beta^{\tau} D^{\tau}_{ikt} \times \texttt{treated}_{ikt} + \epsilon_{ikt} \tag{8}$$

Y_{ikt} is an outcome (such as prescription drug days supply) for patient i, experiencing closure event k, in time period t (months or quarters). The `treated` dummy captures whether patient i is in the control or treated group for closure event k in month t, while D^{τ} is a set of event time dummies. For treated patients, event time is simply months relative to pharmacy closure, with $\tau = 0$ reflecting the month of closure event k. For control patients, event time is months relative to the counterfactual month of closure (as defined previously, this is 12 months prior to the actual

[25]Using closures 12 months later helps address any concerns about seasonality in the effect of exits.

[26]Drug consumption is observed and well-defined only for living patients in Part D, so I exclude patients who die or who exit Part D during the study period. Appendix Figure A4 shows that the effect of closures on mortality and Part D enrollment is a null, though there is a slight pre-trend in the latter case.

[27]Furthermore, only differential changes across the treated and control groups would threaten my empirical findings.

month of closure). I examine one year before and after the closure event, so $D \in [-12, -2], [0, 11]$, with $D = -1$ set as the excluded month. The parameters of interest are the interaction of these two terms, given by β^{τ}. This term captures the effect of pharmacy closures relative to the month before closure. The last term is the econometric error (ϵ_{ikt}).

The stacked design addresses recent concerns about staggered treatment DiD and relies on the standard parallel trends assumption for identification: Drug use patterns of the treated patients would follow the same trend absent the closure as the control patients. A sufficient underlying assumption would be that while pharmacy closures themselves are non-random, the exact timing of closure within a given year is independent of all other factors correlated with patient drug use. While not directly testable, I provide a variety of evidence to support this assumption. First, I show in Table 1 that the characteristics of the treated and matched control patients are similar. Because the control group faces a closure one year in the future, these similarities illustrate that the timing of closures is unlikely to be driven by different trends in patient characteristics or drug demand. Second, Figure 2 illustrates that the pre-trends are parallel, which I show more carefully after implementing the DiD in Figure 3. The lack of pre-trends, combined with a sharp change in the treated group in the month of exit provide further reassurance that the observed effects are unlikely to be driven by different use trajectories of patients using pharmacies that close a year apart.

The primary remaining concern is sudden shocks highly correlated with pharmacy exit. The focus on Medicare patients and the fact that the vast majority of pharmacies are not owned by other providers mitigate concerns about correlated labor or health care market shocks. Furthermore, in Section 4.6, I conduct a robustness check using patients losing access to a pharmacy due to state-wide plan exclusions (which are unlikely to be driven by local shocks) and find qualitatively similar results.

I estimate the model using both linear and Poisson specifications. Because the drug dispensing data contains zeros, the Poisson specifications allow interpreting coefficients as percent changes without requiring the ad hoc adjustments that a logarithmic transform would require. To allow for comparisons across patient and drug groups with different baseline levels of drug use, I present the Poisson specifications, but the results do not differ substantially.[28] This regression is run on 695,421 treated patients, all of whom are randomly matched to exactly one control. While all the treated patients are unique, a given patient can appear in the regression both as a treated and potentially more than once as a control.[29] However, because all patients are matched to exactly one pharmacy exit (the level of treatment), clustering standard errors at the pharmacy exit level corrects for this duplicated data (Deshpande and Li, 2019).

4.5 Estimates of Reduced-Form Effects

As shown in Panel (a) of Figure 3, closures cause a significant 1.2% decline in drug use in the quarter after closure. I bin the monthly estimates to quarters to improve precision and readability and estimate Equation 8 using a Poisson specification with days supply of prescription drugs as the outcome variable measuring drug consumption. The lines shown reflect 95% confidence intervals. Appendix Figure A5 shows that the results are similar using a linear specification; the results from the linear model also match exactly the difference between the treated and control groups relative to month -1 shown in Figure 2. Panel (c) shows the results of the Poisson specification, illustrating that closures cause a 1.2% decline in drug use in the quarter after closure. However, all of these figures show swift attenuation of the decline that reverses sign significantly three quarters after closure.

This overall effect masks substantial heterogeneity across multiple dimensions. First, I consider splits of the population by income status. In particular, I first flag patients as low-income if they receive the Part D low-income subsidy (LIS) for more than half of the year containing the first month of the pre-period. Panel (b) of Figure 3 shows the results of this regression. There is a clear immediate difference in the response to closures by income, with low-income patients experiencing a drop more than twice as large in the quarter following closure. The effect remains negative and significant for low-income patients for twice as long as non-low-income patients. The point estimate of the effect of closure flips sign for both groups by a year after closure, though the long-run increases are noiser and slightly smaller for low-income patients and are thus not statistically

[28] The fact that the pre-period outcomes are not just parallel, but similar in levels, allow for parallel trends in both the linear and Poisson specifications.

[29] Because each treated patient is matched to a control patient who experiences a closure 12 months in the future, in cases where there are fewer control than treated patients for a given closure month, some controls will be duplicated.

significant.

The interpretation of the impact of closures depends on the types of drugs and patients that are affected. I first use classifications from Chandra, Gruber and McKnight (2010) (CGM) to split the drugs into three categories. The authors use a clinician panel to classify drugs based on whether discontinuation is likely to lead to an adverse event within a month, a year, or longer/never. I then consider a subset of chronic cardiac drugs typically thought to be high-value. See Appendix C for details on these classifications. Panel (a) of Figure 4 plots the three categories in the CGM classification while Panel (b) shows the cardiac subset plotted against the effect for all drugs for comparison. Both panels show similar patterns in the treatment effects across the different types of drugs. These chronic cardiac drugs and drugs classified by CGM as likely to prevent adverse events within a month or year are often used in the literature as proxies for "high-value" medications.[30] Under this assumption of ex-ante value, these results show that closures do not primarily reduce consumption of drugs likelier to be low-value. Of course, drug-level definitions of value necessarily ignore patient heterogeneity.

Given the observed heterogeneity in the effects of pharmacy closures across patient income, it is natural to consider whether there is similar heterogeneity at the pharmacy level. Recent policy interest has focused specifically on certain types of closure thought to be especially harmful for patients. For instance, closures of rural or independent pharmacies have been the subject of particular concern, as have closures that create so-called "pharmacy deserts."[31] These types of exits, particularly in rural areas, are the likeliest to increase the distance a patient must travel to the nearest pharmacy. While noisy, Figure 6 shows that the effects of desert-creating pharmacies in rural areas is statistically indistinguishable from other closures.[32] The point estimates are noisy (reflecting the rarity of these types of exits),[33] but I can reject that rural desert-creating closures reduce long-run drug use by more than 0.9% at the 95% level. Similarly, Appendix Figure A6 shows that independent, rural, and the broader set of desert-creating closures (i.e., not just rural)

[30] See for example, Lavetti and Simon (2018) or Choudhry et al. (2011).

[31] For examples from the health policy literature, see Lazaro, Ullrich and Mueller (2022) or Guadamuz et al. (2021). For examples of concern from policy makers, see this release from Senator Katie Britt (R-Alabama) about a congressional hearing on Pharmacy Benefit Managers or the proposed "Pharmacists Fight Back Act."

[32] I define desert-creating exits as closing pharmacies that accounted for more than half of dispensing to patients living in the Zip code of closure in the pre-period. I provide more details in Appendix B.

[33] Desert-creating exits account for 11% of patient-weighted exits. Desert-creating exits in rural areas specifically account for just 4%.

all cause declines in patients' drug use that are no larger than average.

Finally, I use the breadth of the Medicare claims data to estimate the effect of closures on patient outcomes. I focus specifically on hospitalization claims and Medicare hospital spending.[34] Because I do not observe non-Part D health care use for patients in privatized Medicare plans (Medicare Advantage), I subset to patients enrolled in Traditional Medicare. Appendix Figure 5 shows that for both hospital claims and spending, I cannot rule out increases or decreases in any quarter. However, given the modest impacts on drug use, I cannot reject existing estimates in the literature of the effects of drug consumption on hospitalizations.

4.6 Robustness

Even with parallel pre-trends, shocks highly correlated with pharmacy exits still pose an identification threat. To address this possibility, I consider a more subtle type of pharmacy "exit" in which a chain is cut out of the network of a given Part D insurer.[35] While rarer than physical closures, these events occur at the state or nationwide level, and thus are unlikely to be correlated with any local shocks. I define these exclusions empirically using a leave-one-out approach with the claims data.[36] First, for each three-digit Zip code, I calculate total annual days supplied at the state-chain-plan level, excluding the three-digit Zip in question. I then determine for each three-digit Zip the set of state-chain-plan-year leave-one-out tuples where dispensing declines by more than 75% from the prior year, but neither the overall chain nor the overall plan saw declines of more than 20%.[37] These observations comprise the treated group. Although independent pharmacies can also be excluded from insurer networks, the leave-one-out empirical approach relies on the fact that each chain has multiple locations. Thus, these analyses are restricted to exclusions of chain pharmacies.

I define treated patients as those shopping at an excluded pharmacy 7-12 months prior to

[34] Whether patients make rational decisions about drug consumption determines whether negative health outcomes represent a welfare loss. However, because other health care use is heavily subsidized by Medicare and is not borne by the closing pharmacy, closure-induced spending changes are externalities.

[35] I examine full network exclusions and not changes a chain's "preferred" status. Patients typically pay lower copays at their insurers' preferred pharmacies, (Starc and Swanson, 2021), but full network exclusions preclude any insurance benefits at the excluded pharmacy.

[36] CMS publishes lists of pharmacy networks of each Part D plan, but preliminary analyses showed that these lists are insufficiently precise to pinpoint true exclusions.

[37] The 75% threshold helps account for exclusions that do not align exactly with the calendar year and for limited exceptions a plan may make for patients. The final restriction helps prevent counting cases where a chain or plan dramatically shrinks overall as an exclusion.

exclusion and control patients as those on the same plan but who shop at a non-excluded pharmacy 7-12 months prior to the exclusion. Then, I estimate the same event study as in Section 4.5, except that I also show two years of pre-period results. Panel (a) of Appendix Figure A7 shows a clear, sharp decline in prescription drug consumption in patients experiencing a network exclusion relative to the control group. Panel (b) plots these effects on the same axes as the results of chain *closures*. The one difference here is that I specify period -2 instead of period -1 as the excluded level, which helps address the anticipatory spike in period -1 in the exclusion event study. That difference aside, the overall effects of the exclusions and closures are similar. If anything, the effects of the exclusions are slightly larger, ameliorating concerns that correlated local shocks are an identification threat for the closure analyses.

4.7 What We Can and Can't Learn from the Reduced Form

What mechanisms explain the attenuation and eventual reversal in the intial effect of pharmacy closure? Returning to the conceptual model discussed in Section 3, recall that changes in permanent, per-period costs, reallocation, and switching costs effects comprise the overall reduced-form effect. The attenuation in the effect size in the quarters following closure is consistent with pharmacy exits imposing a one-time disruption or switching cost that patients must pay to start using a new pharmacy, capturing the effort and attention required to change to a new pharmacy or to get the prescriber to transfer the prescription. However, permanent increases other price or non-price costs (such as copays or travel distance) should not cause this type of time-varying effect. Furthermore, the lack of larger effects of desert-creating exits and similar attenuation among high and low-income patients suggest that neither changes in distance costs nor prices are the primary mechanism driving the observed results.[38]

Neither of these effects, however, can explain the reversal in the sign of the effect by a year after closure. The model illustrates that reallocations to higher-dispensing pharmacies may be one explanation for the long-run increases, especially if the negative impacts of disruption attenuate over time. The key parameters needed to test this hypothesis are the effects of pharmacies on patient drug consumption, and specifically, whether there is heterogeneity in these effects. Because

[38] Low income patients face low prices irrespective of their choice of pharmacy.

the reduced-form does not provide a way of estimating these impacts, I instead use an AKM-style model.[39]

5 Heterogeneity in Pharmacies' Impact on Patient Drug Use

I estimate pharmacy-level effects on drug consumption and find substantial heterogeneity. Then, using these fixed effects, I show that closing pharmacies are disproportionately low-dispensing and that low-income patients reallocate to higher-dispensing pharmacies after exit. Finally, I leverage the limited costs faced by low-income patients to separately quantify switching costs and reallocation effects from pharmacy exits and find that the overall reduced form effect obscures these mechanisms pushing in opposite directions.

5.1 Adapting Conceptual Model to Examine Pharmacy Fixed Effects

At the cost of adding structure to the patients' drug consumption problem, a two-way fixed effects model (in the style of Abowd, Kramarz and Margolis (1999)) provides a way to isolate pharmacy fixed effects. To start, I return to Equation 3 and group the terms, as shown below.

$$y^*_{it} = \underbrace{\alpha_i + q_j}_{\text{FEs}} + \underbrace{h_{it} - c_{ijt}}_{\text{confounders}} \tag{9}$$

The pharmacy fixed effects (q_j) are the parameters of interest, but regressing drug consumption on the two fixed effects will be biased due to omitted variable bias from the two confounders. In particular, the assumptions that must be satisfied for causal identification are that patients do not systematically sort to certain types of pharmacies based on their health shock (exogenous mobility) and that the patient and pharmacy effects are additively separable (no match effects) (Card, Heining and Kline, 2013).

I thus implement two adjustments to Equation 9 to make these assumptions more plausible. First, instead of a patient fixed effect, I include a patient-by-prescriber-by-insurer fixed effect (α_{ikl}).

[39] I also rule out two alternate explanations outside my model. First, Appendix Figure A8 shows similar patterns in the treatment effects for the subset of drugs that treat acute conditions. This helps rule out concerns about stockpiling or intertemporal substitution, under the assumption that patients are less likely to have stocks of drugs that treat acute conditions. Second, Appendix Figure 5 illustrates no effects on hospitalizations, suggesting that the long-run increases in drug use are not driven by worsening patient health.

Because patient moves across pharmacies identify the q_j term, this three-way interacted fixed effect only counts moves where the patient's doctor and insurer remain unchanged.[40] Thus, α_{ikl} absorbs any time-invariant role the patient's doctor or insurer plays in their prescription drug use. Second, I estimate the model using only low-income patients to avoid match effects stemming from prices. Given the reduced form suggests that travel distance matters little for drug consumption, the fact that low-income patients often face near-zero prices set by Medicare help shut down the cost term.[41] Combining these modifications yields Equation 10, where the unobserved h_{it} and c_{ijt} are combined into an error term: ϵ_{ijt}. Even with these adjustments, the AKM assumptions are still quite strong. I thus present standard robustness checks at the end of Section 5.4.

$$y^*_{it} = \underbrace{\alpha_{ikl} + q_j}_{\text{FEs}} + \epsilon_{it} \tag{10}$$

To provide a rough benchmark for the estimated pharmacy effects, I estimate a similar regression to isolate the effects of prescribers. By replacing the patient-by-prescriber-by-insurer fixed effect with a patient-by-*pharmacy*-by-insurer fixed effect, I can isolate a fixed effect for prescribers, as shown below in Equation 11.

$$y^*_{it} = \underbrace{\alpha_{ijl} + q_k}_{\text{FEs}} + \epsilon_{it} \tag{11}$$

By relying on patients switching prescribers while their pharmacy and insurer remain the same, this regression is identified from an entirely separate set of patient moves. While this makes direct comparisons to the pharmacy effect from Equation 10 more challenging, I use this comparison merely to provide some additional context for the magnitude of the pharmacy effects I estimate.

5.2 Sample Selection

Starting from the 20% Part D patient sample, I make two major sample restrictions to implement the two-way fixed effects model. First, I keep only patient-years where the patient is enrolled in the Part D low-income subsidy and alive for the entire year. Patients are then matched to their modal

[40] Otherwise a new fixed effect is assigned, essentially treating the observation as a new "patient."

[41] Non-low-income patients face different prices at different pharmacies based on their insurer's network.

prescriber, insurer, and pharmacy for each calendar year. The prescriber and dispensing pharmacy are both directly rcorded on each Part D claim and I define the insurer as the contract ID listed on the claim.[42] Log total annual days supply of drugs is the baseline outcome of interest used to define modal providers. I then drop patient-years where either the modal prescriber or pharmacy does not have an NPI (to avoid conflating changing ID formats with changing providers) or where the patient does not consume any prescription drugs.[43]

Because one major goal of estimating this model is to incorporate the AKM pharmacy fixed effects into the reduced-form analyses, I impose further restrictions to avoid a mechanical relationship between the two. For any patient treated by a pharmacy closure, I drop all patient-years greater than or equal to the year prior to the closure. This restriction ensures that none of the variation used to estimate the reduced-form is used in estimating the model.

5.3 Variance Correction, Summary Statistics, and Tests of Validity

With the sample defined as described above, I now turn to estimating the AKM model. So long as the identification assumptions described previously hold, the modeled pharmacy fixed effects are a consistent estimate of the effect of each pharmacy on patients' drug consumption. However, the impact of pharmacies overall is captured by the variance in these pharmacy effects. Unlike the fixed effects themselves, the estimated variance is biased upwards due to noise in the estimated fixed effects. This widely known issue arises from estimating fixed effects from a finite (and often very small) number of movers and is thus often referred to as "limited mobility bias" (Kline, Saggio and Sølvsten, 2020; Bonhomme et al., 2023). I implement the correction proposed by Kline, Saggio and Sølvsten (2020) that addresses this bias by restricting the sample to the largest leave-one-out connected set and then implements a heteroskedastic-robust bias correction on this subsample.[44] Restricting to the leave-one-out connected set requires dropping 12% of pharmacies

[42]This definition of "insurer" reflects a balance between providing as much specificity about the patients' plan as possible and also maintaining a large connected set to estimate the model. Each contract ID reflects a set of Medicare Advantage or standalone Part D plans offered by a single insurance company. Thus, this variable will capture switches across insurance companies or between Medicare Advantage and standalone Part D plans, but not necessarily smaller changes across plans offered by a single insurer.

[43]In years with no consumption, I do not observe the patients' pharmacy or prescriber. However, the overall sickness of the Medicare population means that zeros are rare.

[44]The largest connected set of a network is the largest subset of nodes that are connected by at least one chain of edges. The largest leave-one-out connected set is the largest connected set where all links are maintained even if any given patient-by-prescriber-by-insurer is dropped from the sample.

and 3% of patient-by-prescriber-by-insurer observations, leaving 66,631 pharmacies and 4,017,993 unique patient-prescriber-insurer combinations in the leave-one-out set. Appendix Table A1 provides additional summary statistics for both this regression as well as the prescriber version.

After estimating the fixed effects on the leave-one-out connected set, I probe the validity of the AKM assumptions. Panel (a) of Appendix Figure A9 shows a balanced event study of log annual drug consumption in the year of a move and two years prior for pharmacy movers.[45] The figure bins movers based on the quartile of the fixed effect of the origin and destination pharmacy and focuses on patients who start at a pharmacy in either the lowest or highest quartile. Moves between pharmacies in the same quartile of fixed effects have little impact on patients, but moves to higher or lower-dispensing pharmacies have opposite and largely equal in magnitude effects on log prescription drug consumption. The pre-trends in the year prior to the move are relatively similar across the eight groups, with some deviations in the 4 to 2 group and 4 to 1 group. However, these differences in the pre-trends are small compared to the change due to the move and affect a relatively small set of movers. Panel (b) splits the patient-prescriber-insurer and pharmacy fixed effects in quintiles and plots the average residual of Equation 10 for observations within each bin. Large or systematic patterns in the residuals can suggest the presence of match effects. Instead, I find relatively small residuals (most are less than 0.3% and the largest is 2.3%), with the largest residuals in the first and fifth quintiles.

Appendix Figure A10 shows the same event studies and residuals plot for the prescriber fixed effects. Relative to pharmacy movers, The differences in pre-trends in the event study (Panel (a)) and the residuals (Panel (b)) are larger, suggesting stronger match effects. Because match effects typically cause overestimates of the pharmacy/prescriber effects (Bonhomme, Lamadon and Manresa, 2019), I take this as evidence that if anything, my estimates of pharmacy effects relative to prescriber effects may understate the relative impact of pharmacies.

[45]The three-way patient-by-prescriber-by-insurer fixed effects cause the panels to be short. Thus, even the restriction in this figure of subsetting to movers with data for the two years prior to the move already requires dropping 69% of movers, which is why I do not show longer event studies.

5.4 Variance of Pharmacy Fixed Effects

Figure 7 shows the share of variance explained after implementing the Kline, Saggio and Sølvsten (2020) correction.[46] The first blue bar in Panel (a) shows that pharmacies (using Equation 10) explain 1.2% of the overall variance in prescription drug consumption. The latter blue bar illustrates sorting of high-use patients to high-dispensing pharmacies; the covariance of the pharmacy and patient term is positive and explains 1.7% of the total variance. Panel (b) shows the raw variances, illustrating that the variance in the pharmacy effect is 0.009. The red bars in both panels show the effects of the separate prescriber regression (from estimating Equation 11). The variance in the prescriber fixed effect is 0.021, accounting for 2.3% of the total variance in the prescriber model. The effect of pharmacies is thus roughly half that of prescribers—50% based on the share of variance explained in their respective models, or 44% based on the variance itself.[47]

While pharmacies' influence over patients' drug consumption has received limited prior attention, there is much more evidence that prescribers matter (Zolnierek and DiMatteo, 2009; Laird and Nielsen, 2016). Thus, despite the two models being estimated on different sets of patient moves, the estimated prescriber effects provide a valuable benchmark for comparison. I interpret the fact that pharmacies have roughly half the effect of prescribers as evidence that pharmacies play a meaningful role in patients' drug consumption. Appendix Table A2 shows that higher pharmacy fixed effects are positively correlated with mail-order pharmacies and negatively associated with chain pharmacies and exiting pharmacies.[48]

5.5 Combining AKM Estimates With Closure Analyses to Test Mechanisms

Using the pharmacy fixed effects estimated in the AKM model, I can now test the hypothesis that reallocations to higher-dispensing pharmacies drive the long-run increases in drug consumption following pharmacy exit. To match the population used to estimate the AKM model, I focus on low-income patients.I also drop any patients from these reduced-form analyses who are included

[46] All of these analyses are implemented using the `PyTwoWay` package developed by Bonhomme et al. (2023).

[47] Panel (b) of Figure 7 shows the raw variance estimates of all the parameters across the two specifications, illustrating the broadly similar overall variance and share explained by the other terms.

[48] Note that this fact does not necessarily imply that closures cause patients to move to higher-dispensing pharmacies. For instance, if exiting pharmacies are above the average dispensing of their neighbors, closures could cause reallocations to lower dispensing pharmacies, even if the closing pharmacies are below the overall average.

in the AKM model,[49] and subset to the set of patients shopping at a closing pharmacy for which I can estimate a fixed effect in the AKM model, which requires dropping 43% of observations in total.[50] This restriction is by necessity, but Appendix Figure A11 shows no meaningful differences in the effect of closures on the overall set of closures and the subset with an estimated fixed effect.

After imposing the sample restrictions described above, Panel (a) of Figure 8 shows the same event study estimated previously, but with the pharmacy fixed effects as the outcome.[51] While there are no changes prior to exit, the average pharmacy fixed effect of pharmacies visited by a low-income patient increases significantly right after closure. The magnitude is about 0.01 log points, indicating that patients move, on average, to pharmacies that increase patients' drug use by 1%.[52]

Now that we have the average change in the pharmacy dispensing effect, we can use the conceptual model presented in Equation 7 to separate out the components comprising the overall closure effect. Recall that there are three components: permanent changes in costs, changes in the pharmacy effect (reallocation), and temporary switching costs. I showed previously using the desert-creating closures that distance does not appear to drive drug consumption decisions, which combined with limited changes in prices for low-income patients mean that the change in costs is negligible.[53] Thus, the only remaining pieces are the reallocation and disruption effects and by subtracting the reallocation from the overall effect, we can recover the effect of disruption.[54] Panel (b) of Figure 8 shows this decomposition. In the short-run the disruption effect dominates, but as this effect attenuates, the reallocation effect causes the sign on the overall effect to reverse.

[49] Note that I already dropped all patient-years in the year prior to closure and later for patients in the closure reduced-form analyses from the AKM. This extra restriction removes any patients who were used to identify pharmacy fixed effects in one of the earlier years before the closure. Thus, after imposing this restriction, there is truly no overlap in the set of patients used to identify the AKM pharmacy fixed effects and the analyses shown in this subsection. However, this restriction makes little difference in the results.

[50] Dropping patients whose closing pharmacy does not have a fixed effect costs 33% of the original sample, while removing any patients from the reduced form who were used to identify any firm effect in the AKM model requires dropping another 10% of the original sample.

[51] For patients that fill prescriptions at more than one pharmacy in a given quarter, the fixed effect is the quantity-weighted average across the pharmacies they shop at. Note that the panel is unbalanced because the fixed effect is undefined for patients without any consumption in a quarter, so these observations are dropped.

[52] Overall, about 57% of patient-weighted exits are below the median pharmacy dispensing level.

[53] There are some slight, but generally insignificant decreases in prices charged to low-income patients following closure. Using a broad range of price elasticity estimates (e.g., Hollrah (2024)), these small price changes never substantially impact the results. Furthermore, there is some debate as to whether low-income patients actually pay these copays at all (Gross, Layton and Prinz, 2022).

[54] Estimating the AKM model in logs and the reduced-form using a Poisson specification mean that the effect sizes can be interpreted as percent changes in both cases, allows me to directly compare these estimates.

To further validate the causal effects of the pharmacy fixed effects, I show the effects of exits of above versus below-median dispensing pharmacies in Panel (a) of Figure 9. I estimate the event study (Equation 8) separately for each of these two groups. Importantly, I use the same split for the control pharmacies to ensure that patients experiencing a high-dispensing (low-dispensing) pharmacy closure today are compared to patients experiencing a high-dispensing (low-dispensing) closure a year from now. Reassuringly, all of the long-run increases are driven by low-dispensing exits. Pre-trends are relatively flat in both cases and patients in both groups experience a similar immediate decline in the quarter of exit, but below-median dispensing exits cause drug use to increase by 1% in the third quarter after closure. On the other hand, high-dispensing exits have slightly negative, but imprecise, effects on patients' drug use even a year after exit.

Finally, I also condition on the fixed effect of patients' post-period pharmacy to explicitly study reallocations from low to high-dispensing pharmacies and vice-versa. These regressions must be interpreted with caution because they condition on an outcome variable in the treated group and exclude any patients that do not take drugs on the extensive margin. With these caveats in mind, Panel (b) of Figure 9 shows that the long-run increases are specifically driven by reallocations from low to high-dispensing pharmacies. None of the three other types of moves (low to low, high to high and high to low) cause long-run increases.[55]

Note that while the AKM model relies on strong assumptions, the implications of violations of these assumptions are different in the reduced-form analyses. In particular, if the exogenous mobility assumption is violated by patients sorting to certain types of pharmacies (e.g., mail-order) after receiving a negative health shock, then the AKM model will *overestimate* the impact of pharmacies. However, my claim that exits of low-dispensing pharmacies drive long-run increases in drug consumption only uses the AKM fixed effects to classify pharmacies as high or low-dispensing. Thus, any noise in the fixed effects or biases from these types of violations of the AKM assumptions will potentially cause me to misclassify some low-dispensing exits as high-dispensing and vice-versa, causing me to *underestimate* the differences between low and to high-dispensing exits.

So far, I have purposely avoided normative statements about whether higher-dispensing phar-

[55]Though note that there is a bit of an asymmetry because moves from high to low-dispensing pharmacies do not appear to cause larger long-run declines.

macies are higher quality. More dispensing can be welfare-improving if it increases consumption of high-value medications that patients otherwise underconsume due to prices above marginal cost, behavioral biases, poor information, or externalities. Understanding pharmacies' role in patient consumption decisions is important regardless, but I also provide evidence that dispensing of socially harmful drugs (opioids) do not drive the pharmacy effects. In particular, I estimate the correlation between the the share of a pharmacy's dispensing that is opioids and their model-estimated fixed effect. I find that these are slightly negatively correlated (-0.03) illustrating that opioids are not driving the model estimates.

5.6 Policy Implications

Increasing pharmacy reimbursements is a growing policy priority, often with the stated goal of preventing pharmacy exits and promoting drug access. However, given the heterogeneity I document in pharmacies' effects on patient drug consumption, the effects of these interventions depend on the types of exits prevented. The magnitudes I estimate also help compare the cost-effectiveness of subsidizing prescription drugs through higher pharmacy reimbursements versus other policies the government could pursue.

To consider these trade-offs I examine a recent increase in pharmacy reimbursements in the California Medicaid program. This policy effectively added a 30% increase to dispensing fees paid to small pharmacies beginning in 2022. Using non-chain pharmacies as a proxy for small ones,[56] I find suggestive evidence that this policy was effective in preventing exits, but that the policy likely cost millions of dollars per averted closure (details in Appendix D).

While this policy likely prevented short-term declines in drug use following each of the averted exits, the long-run effect depends on the types of closures prevented. Under the assumption that the policy prevents the average closure, it would *reduce* long-run drug use by 1.1%. Any policy-driven increases in drug use would need to be large to make this policy cost-effective.

[56] I do not have claims data for these years, so I rely on existing work showing that non-chain retail pharmacies tend to be much smaller than chains (Ladsariya et al., 2023).

6 Conclusion

The fact that patients often do not take high-value medications, even when free, suggests a potential role for other non-price costs in driving patients' medication use. This paper studies a set of non-price mechanisms related to a key node in the drug supply chain: pharmacies. Recent policy concerns have largely focused on the impacts of pharmacy exits, but there has been little causal evidence of whether closures, and pharmacies more generally, matter. The reduced-form results illustrate that closures cause initial declines, but long-run increases in prescription drug consumption. My conceptual model illustrates that closure-driven reallocations from lower to higher-dispensing pharmaciesa could be one explanation for the long-run effect. I develop and estimate an AKM-style model on low-income patients to test this hypothesis and examine the broader impacts of pharmacies. I find significant heterogeneity in pharmacies' impacts on patient drug use equal to roughly half the variation between prescribers. However, closing pharmacies are disproportionately lower-dispensing and reallocations to higher-dispensing pharmacies drive all the long-run increases in low-income patients' drug use. These findings illustrate the risks of broad efforts to prevent closures as a way to promote access to prescription drugs. Not only are these policies often costly, but they may be counterproductive by preventing exits of low-dispensing pharmacies. However, my results also illustrate the potential power of targeted policies that move patients to higher-dispensing establishments or that encourage exits of low-dispensing pharmacies.

References

Abowd, John M, Francis Kramarz, and David N Margolis. 1999. "High wage workers and high wage firms." *Econometrica*, 67(2): 251–333.

Agafiev Macambira, Danil, Michael Geruso, Anthony Lollo, Chima D. Ndumele, and Jacob Wallace. 2025. "The Private Provision of Public Services: Evidence from Random Assignment in Medicaid." *Working Paper.*

Anderson, Kelly E, Maitreyi Sahu, Michael J DiStefano, Carl V Asche, and T Joseph Mattingly. 2024. "Pharmacy Closures and Anticonvulsant Medication Prescription Fills." *JAMA*, 332(21): 1847–1849.

Badinski, Ivan, Amy Finkelstein, Matthew Gentzkow, and Peter Hull. 2023. "GEOGRAPHIC VARIATION IN HEALTHCARE UTILIZATION: THE ROLE OF PHYSICIANS." *NBER Working Paper.*

Baker, Andrew C, David F Larcker, and Charles CY Wang. 2022. "How much should we trust staggered difference-in-differences estimates?" *Journal of Financial Economics*, 144(2): 370–395.

Bendicksen, Liam, and Aaron S Kesselheim. 2022. "Anticipated efficiencies, real costs: Medicaid managed care organizations and the pharmacy benefit." *Journal of Managed Care & Specialty Pharmacy*, 28(3): 354–361.

Berenbrok, Lucas A, Nico Gabriel, Kim C Coley, and Inmaculada Hernandez. 2020. "Evaluation of frequency of encounters with primary care physicians vs visits to community pharmacies among Medicare beneficiaries." *JAMA Network Open*, 3(7): e209132–e209132.

Bonhomme, Stéphane, Kerstin Holzheu, Thibaut Lamadon, Elena Manresa, Magne Mogstad, and Bradley Setzler. 2023. "How much should we trust estimates of firm effects and worker sorting?" *Journal of Labor Economics*, 41(2): 291–322.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa. 2019. "A distributional framework for matched employer employee data." *Econometrica*, 87(3): 699–739.

Brot-Goldberg, Zarek, Timothy Layton, Boris Vabson, and Adelina Yanyue Wang. 2023. "The behavioral foundations of default effects: theory and evidence from Medicare Part D." *American Economic Review*, 113(10): 2718–2758.

Brot, Zarek, Samantha Burn, Timothy Layton, and Boris Vabson. 2024. "Rationing Medicine Through Bureaucracy: Authorization Restrictions in Medicare." *Working Paper.*

Buxbaum, Jason D, Michael E Chernew, A Mark Fendrick, and David M Cutler. 2020. "Contributions Of Public Health, Pharmaceuticals, And Other Medical Care To US Life Expectancy Changes, 1990-2015: Study examines the conditions most responsible for changing US life expectancy and how public health, pharmaceuticals, other medical care, and other factors may have contributed to the changes." *Health Affairs*, 39(9): 1546–1556.

Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace heterogeneity and the rise of West German wage inequality." *The Quarterly journal of economics*, 128(3): 967–1015.

Carroll, Caitlin. 2023. "Impeding Access or Promoting Efficiency? Effects of Rural Hospital Closure on the Cost and Quality of Care." *Working Paper.*

Chandra, Amitabh, Evan Flack, and Ziad Obermeyer. 2024. "The health costs of cost-sharing." *The Quarterly Journal of Economics*, qjae015.

Chandra, Amitabh, Jonathan Gruber, and Robin McKnight. 2010. "Patient cost-sharing and hospitalization offsets in the elderly." *American Economic Review*, 100(1): 193–213.

Chevalier, Judith A, Jason L Schwartz, Yihua Su, and Kevin R Williams. 2022. "JUE insight: distributional impacts of retail vaccine availability." *Journal of urban economics*, 127: 103382.

Choudhry, Niteesh K, Jerry Avorn, Robert J Glynn, Elliott M Antman, Sebastian Schneeweiss, Michele Toscano, Lonny Reisman, Joaquim Fernandes, Claire Spettell, Joy L Lee, et al. 2011. "Full coverage for preventive medications after myocardial infarction." *New England Journal of Medicine*, 365(22): 2088–2097.

Cutler, Rachelle Louise, Fernando Fernandez-Llimos, Michael Frommer, Charlie Benrimoj, and Victoria Garcia-Cardenas. 2018. "Economic impact of medication non-adherence by disease groups: a systematic review." *BMJ open*, 8(1): e016982.

Deshpande, Manasi, and Yue Li. 2019. "Who is screened out? Application costs and the targeting of disability programs." *American Economic Journal: Economic Policy*, 11(4): 213–48.

Dranove, David, Christopher Ody, and Amanda Starc. 2021. "A dose of managed care: Controlling drug spending in medicaid." *American Economic Journal: Applied Economics*, 13(1): 170–197.

Einav, Liran, Amy Finkelstein, and Maria Polyakova. 2018. "Private provision of social insurance: drug-specific price elasticities and cost sharing in Medicare Part D." *American Economic Journal: Economic Policy*, 10(3): 122–53.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams. 2016. "Sources of geographic variation in health care: Evidence from patient migration." *The quarterly journal of economics*, 131(4): 1681–1726.

Fischer, Stephanie J., Heather Royer, and Corey D. White. 2023. "Health Care Centralization: The Health Impacts of Obstetric Unit Closures in the US." *NBER Working Paper*.

Gannaway, Grant. 2019. "Pharmacy Deserts and Medication Adherence." *Unpublished Manuscript*.

Garthwaite, Craig, Tal Gross, and Matthew J Notowidigdo. 2018. "Hospitals as insurers of last resort." *American Economic Journal: Applied Economics*, 10(1): 1–39.

Gross, Tal, Timothy J Layton, and Daniel Prinz. 2022. "The liquidity sensitivity of healthcare consumption: Evidence from social security payments." *American Economic Review: Insights*, 4(2): 175–190.

Guadamuz, Jenny S, G Caleb Alexander, Shannon N Zenk, and Dima M Qato. 2020. "Assessment of pharmacy closures in the United States from 2009 through 2015." *JAMA internal medicine*, 180(1): 157–160.

Guadamuz, Jenny S, Jocelyn R Wilder, Morgane C Mouslim, Shannon N Zenk, G Caleb Alexander, and Dima Mazen Qato. 2021. "Fewer pharmacies in Black and Hispanic/Latino neighborhoods compared with White or diverse neighborhoods, 2007–15: study examines pharmacy "deserts" in Black and Hispanic/Latino neighborhoods compared with White or diverse neighborhoods." *Health Affairs*, 40(5): 802–811.

Heller, David J, Pamela G Coxson, Joanne Penko, Mark J Pletcher, Lee Goldman, Michelle C Odden, Dhruv S Kazi, and Kirsten Bibbins-Domingo. 2017. "Evaluating the impact and cost-effectiveness of statin use guidelines for primary prevention of coronary heart disease and stroke." *Circulation*, 136(12): 1087–1098.

Hollrah, Christopher. 2024. "Prioritizing Value? How Low-Income Patients Respond to Prescription Cost-Sharing." *Working Paper.*

Janssen, Aljoscha, and Xuan Zhang. 2023. "Retail pharmacies and drug diversion during the opioid epidemic." *American Economic Review*, 113(1): 1–33.

Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten. 2020. "Leave-out estimation of variance components." *Econometrica*, 88(5): 1859–1898.

Ladsariya, A, A McLeod, BJ Tevelow, G Noh, and N Sahni. 2023. "Meeting Changing Consumer Needs: The US Retail Pharmacy of the Future."

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

Lavetti, Kurt, and Kosali Simon. 2018. "Strategic formulary design in Medicare Part D plans." *American Economic Journal: Economic Policy*, 10(3): 154–192.

Layton, Timothy J, Nicole Maestas, Daniel Prinz, and Boris Vabson. 2022. "Health Care Rationing in Public Insurance Programs: Evidence from Medicaid." *American Economic Journal: Economic Policy*, 14(4): 397–431.

Lazaro, Edmer, Fred Ullrich, and Keith Mueller. 2022. "Update on Rural Independently Owned Pharmacy Closures in the United States, 2003-2021." *RUPRI Center for Rural Health Policy Analysis www.banko Rural Policy Brief.*

Mourot, Pauline. 2025. "Should Top Surgeons Practice at Top Hospitals? Sorting and Complementarities in Healthcare." *Working Paper.*

Olenski, Andrew. 2022. "Reallocation and the (In) efficiency of Exit in the US Nursing Home Industry." *Working Paper.*

Osterberg, Lars, and Terrence Blaschke. 2005. "Adherence to medication." *New England journal of medicine*, 353(5): 487–497.

Qato, Dima M, G Caleb Alexander, Apurba Chakraborty, Jenny S Guadamuz, and John W Jackson. 2019. "Association between pharmacy closures and adherence to cardiovascular medications among older US adults." *JAMA network open*, 2(4): e192606–e192606.

Sabety, Adrienne. Forthcoming. "The Value of Relationships in Healthcare." *Journal of Public Economics.*

Starc, Amanda, and Ashley Swanson. 2021. "Preferred pharmacy networks and drug costs." *American Economic Journal: Economic Policy*, 13(3): 406–46.

Valliant, Samantha N, Sabree C Burbage, Shweta Pathak, and Benjamin Y Urick. 2022. "Pharmacists as accessible health care providers: quantifying the opportunity." *Journal of Managed Care & Specialty Pharmacy*, 28(1): 85–90.

Zolnierek, Kelly B Haskard, and M Robin DiMatteo. 2009. "Physician communication and patient adherence to treatment: a meta-analysis." *Medical care*, 47(8): 826–834.

7 Tables and Figures

Variable	Control	Treated
Obs.	695421	695421
Days Supply	144.2	144.4
Total Copayment	36.4	36.2
Low-Income	0.412	0.407
Birth Year	1943.55	1943.42
White	0.754	0.751
Disability Qual.	0.335	0.332
Man	0.394	0.392
Desert-Creating	0.109	0.104
Urban	0.784	0.795

Table 1: Covariate Balance. Each treated patient is matched to exactly one control patient (who experiences a closure one year in the future), with replacement. Thus, the number of observations are mechanically the same. See the text for additional details on how the treated and control patients are defined. Days supply and total copayment are at the month level and are averaged over the six-month window used to define the treated and control groups (i.e., months -12 to -7 in event time). Low income refers to the fraction of patients receiving the low-income subsidy. The bottom five variables are all fractions for five binary variables: including race (white versus non-white), whether the patient originally qualified for Medicare due to a disability, sex, whether the pharmacy closure created a "desert" (defined as a closure that accounted for more than 50% of dispensing in the Zip code, and whether the patients' home is in an urban/suburban or rural area (defined at the Zip code level using the Census/USDA RUCA classification, see Appendix B for details.

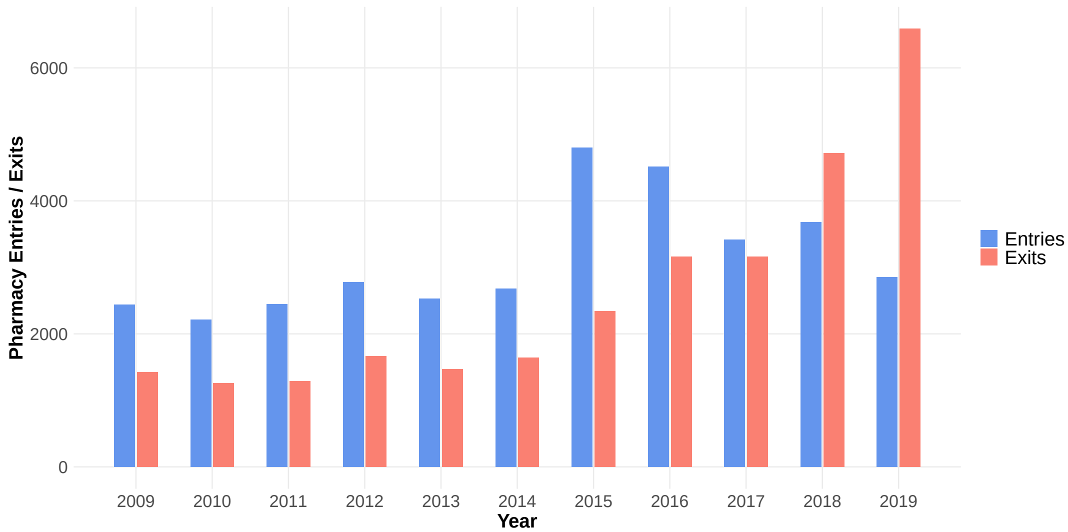

(a) Annual Exits and Entries of Pharmacy Locations

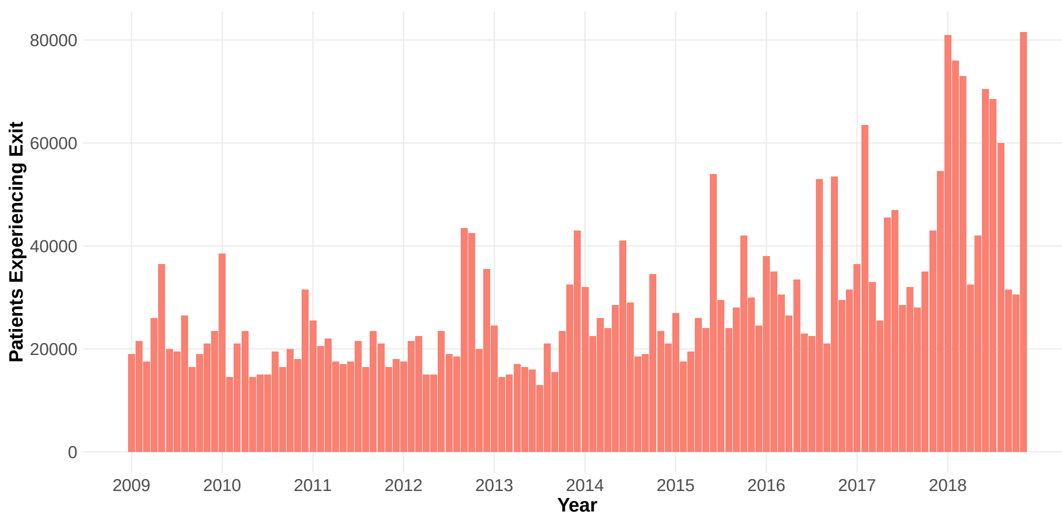

(b) Number of Part D Patients Impacted by Pharmacy Exits Per Month

Figure 1: Panel (a) shows the annual number of entries and exits of pharmacy locations over the study period, excluding pharmacy moves and changes in ownership. Panel (b) describes the number of Medicare Part D patients treated by a pharmacy closure each month. Treatment is defined using the procedure used in the empirical analyses (Section 4.4), restricting to patients shopping at the closing pharmacy 6-12 months prior to closure to avoid selection issues with the patients who shop at a pharmacy right before exit and requiring that the patient be alive and enrolled in Part D for the 24 months surrounding exit. The counts in this figure are multiplied by five to provide the counts for the overall Part D population, rather than just the 20% sample I have.

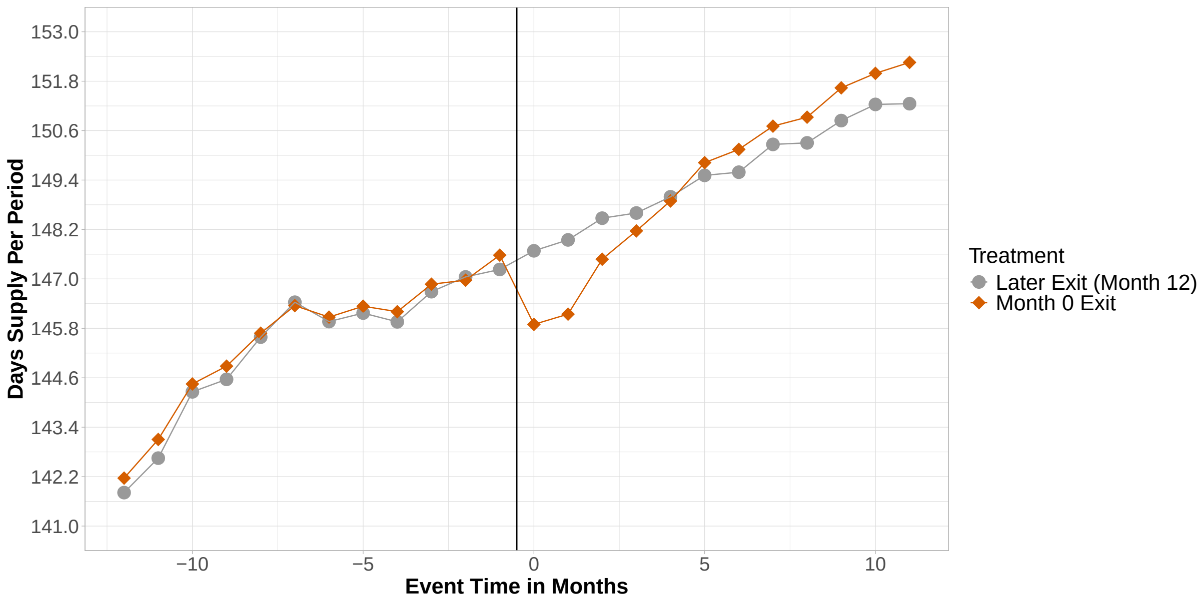

Figure 2: Raw monthly prescription drug days supplied (adjusted for multi-month supplies as described in the main text) around pharmacy closure for treated and control patients. The set of points from -12 to -7 represent the period over which the sample population is defined by those with at least one fill at the closing pharmacy (treated, in orange) or future-closing pharmacy (control, in gray). After this condition is relaxed, the remainder of time periods from -6 to 11 follow this panel of patients. Month 0 is the month of closure.

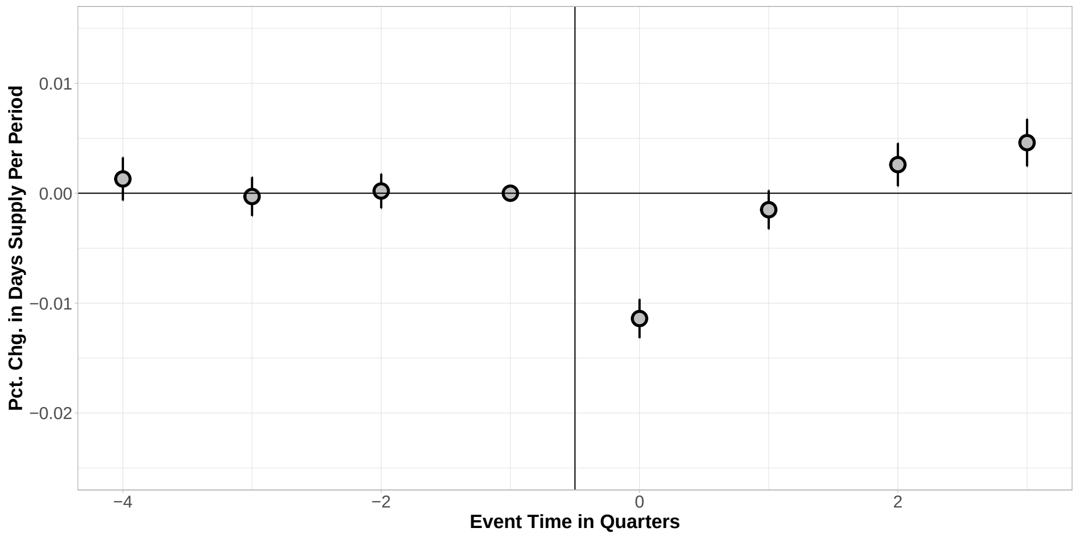

(a) All Patients, Poisson Regression, Months Binned into Quarters

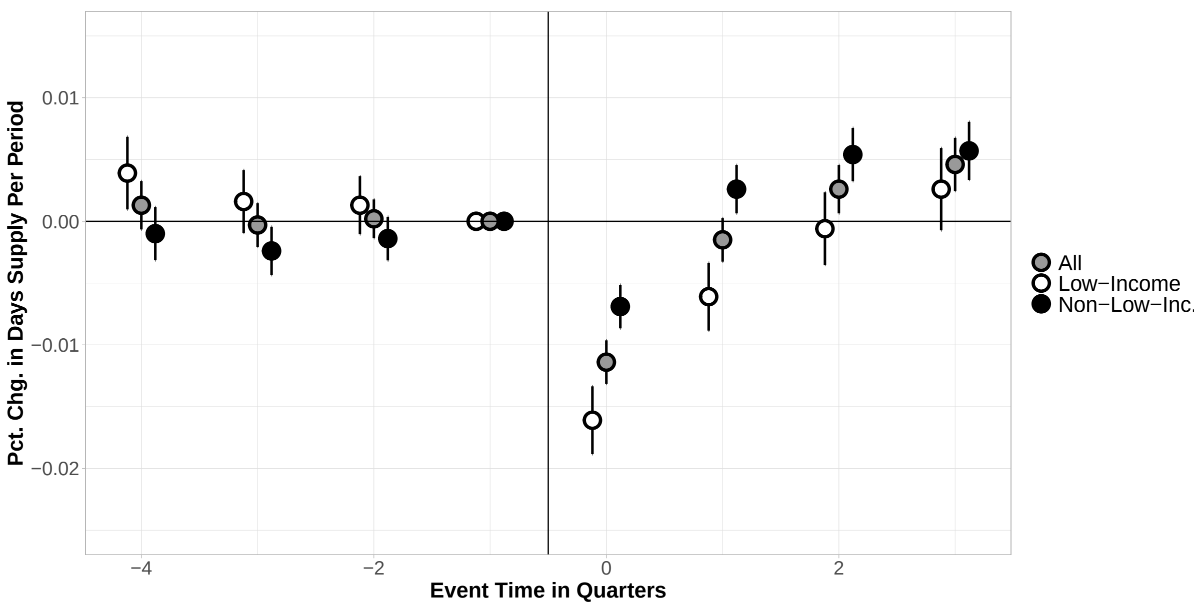

(b) Patients Split by Low-Income Status

Figure 3: Event study coefficients and 95% confidence intervals from Equation 8. The outcome shown here is monthly prescription drug days supply (adjusted for multi-month supplies as described in the main text). Both panels show the estimates from a Poisson specification, with months binned into event-time quarters to increase precision. but all outcomes are still kept at the monthly level. Panel (a) shows the results for all patients, while Panel (b) shows splits by whether the patient receives low-income subsidies in the pre-period. Month 0 is the month of closure. To improve readability, the plot shows the coefficients staggered slightly as a visual aid.

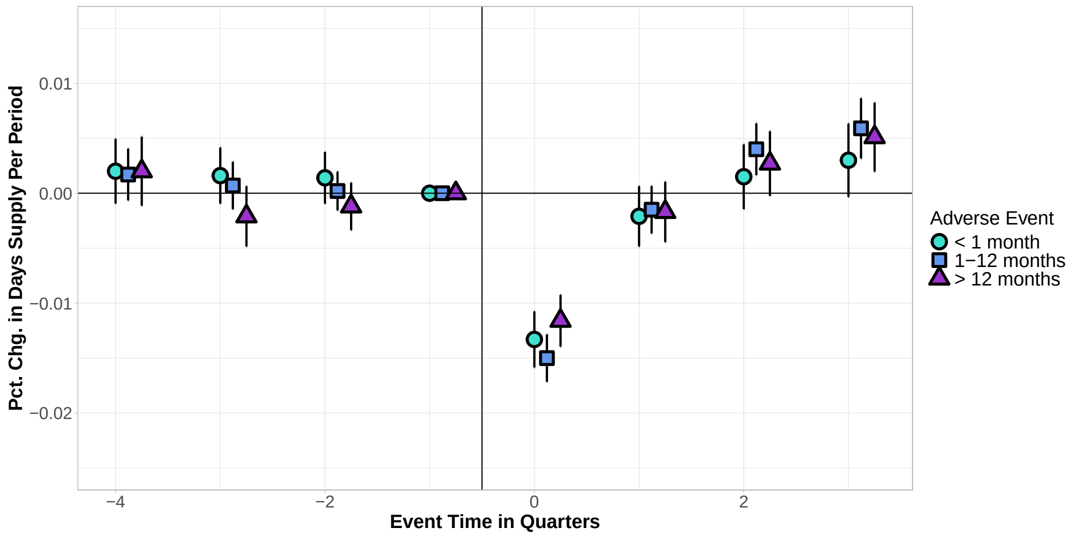

(a) CGM Drug Classes, All Patients

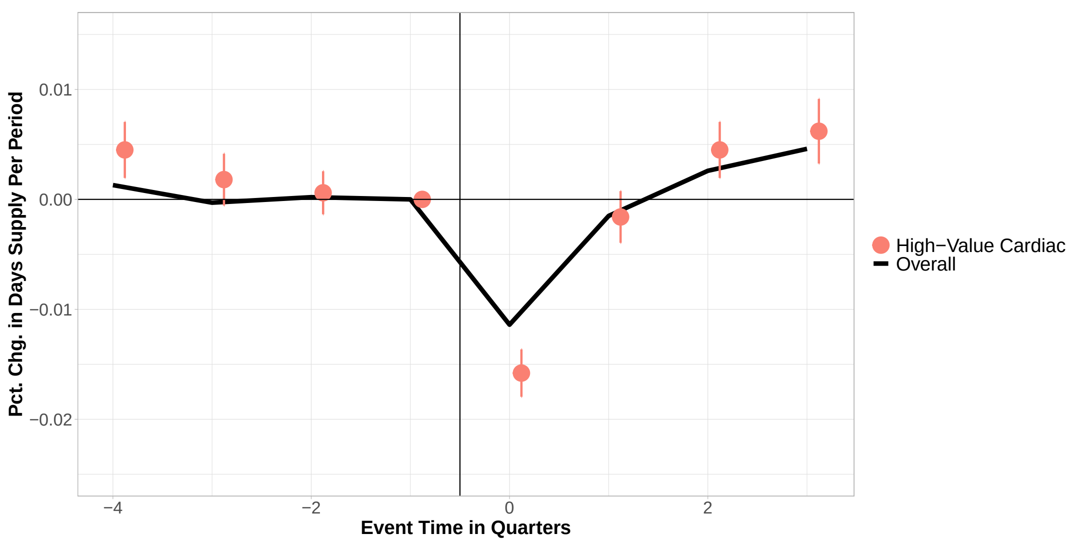

(b) "High-Value" Cardiac Drugs, All Patients (Black Line is Avg. Across All Drugs)

Figure 4: Event study coefficients and 95% confidence intervals from Equation 8. The outcome shown here is monthly prescription drug days supplied split by drug type (adjusted for multi-month supplies as described in the main text). All outcomes are still kept at the monthly level to enable easy comparisons to monthly results). Panel (a) shows outcomes split using the Chandra, Gruber and McKnight (2010) drug classifications, while Panel (b) shows the effect for a specific set of cardiovascular drugs typically thought to be high-value (ACE inhibitors, ARBs, statins, and beta-blockers). Equation 8 is estimated with a Poisson specification; thus, the vertical axis represents percent changes in days supply per month from the reference quarter -1. Period 0 is the period of closure.

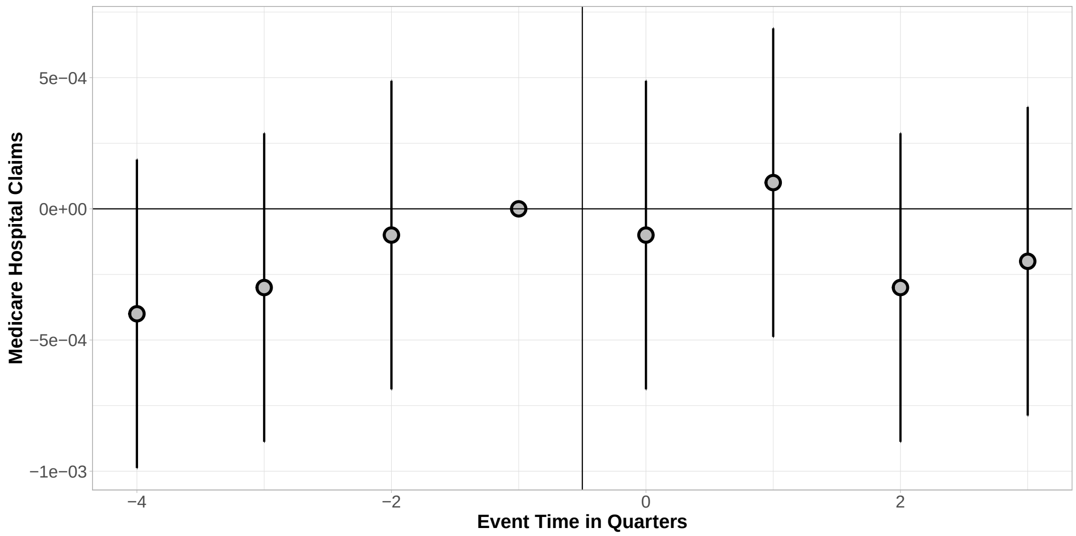

(a) Binary for Hospitalization Per Month

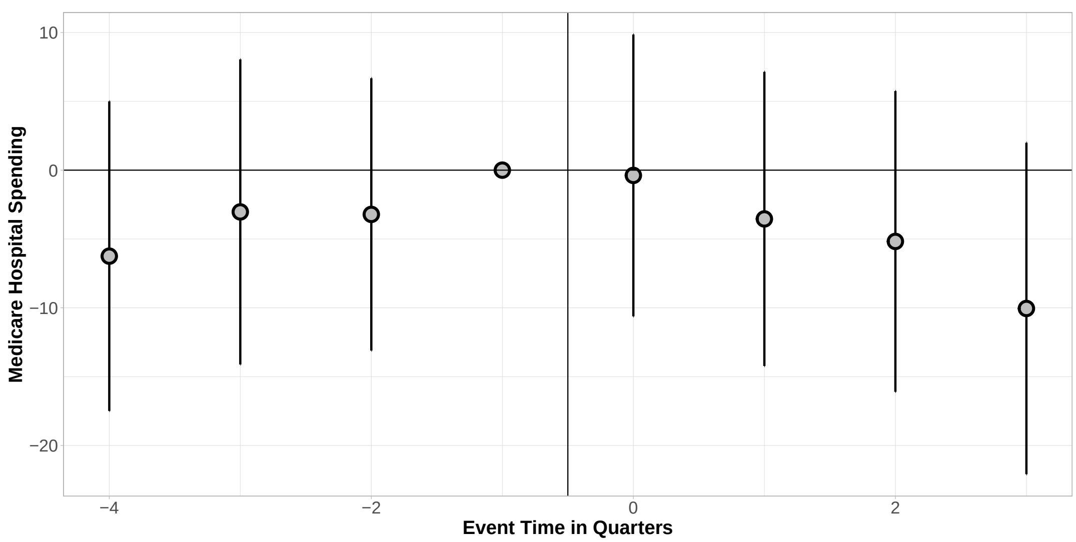

(b) Medicare Hospitalization Spending Per Month

Figure 5: Event study coefficients and 95% confidence intervals from Equation 8. The outcome shown here is a binary for whether the patient has a hospitalization in each month (Panel (a)) and spending by Medicare on hospitalizations (Panel (b)), both binned at quarter level to increase precision. Only Traditional Medicare patients are included and period 0 is the period of closure.

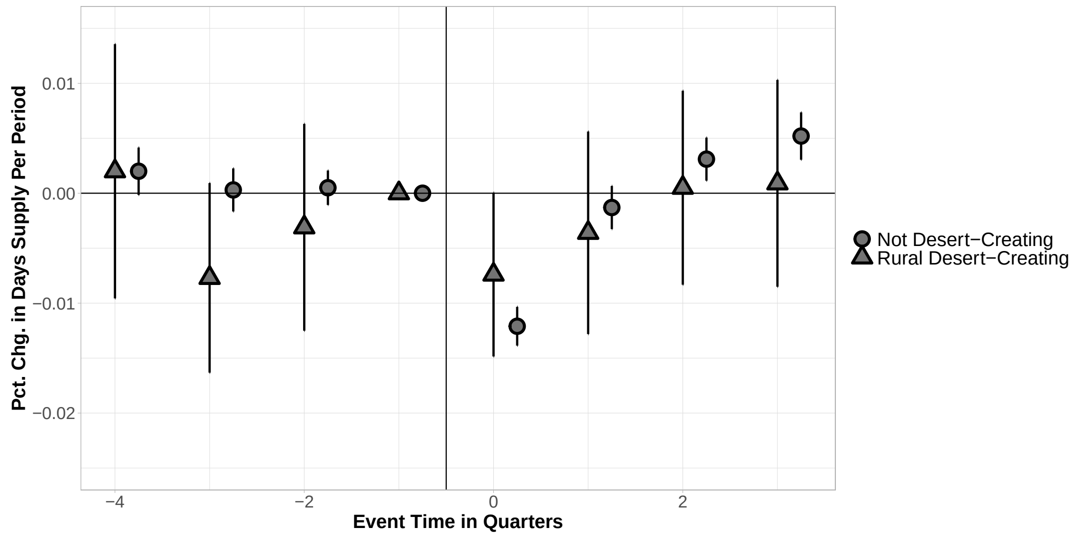

Figure 6: Event study coefficients and 95% confidence intervals from Equation 8 using the Poisson specification and splitting patients based on characteristics of their closing pharmacy. In particular, the "Rural Desert-Creating" estimates restrict to the 4% of closures defined as desert-creating pharmacies in rural Zip codes as defined in the text. The other series plotted excludes desert-creating closures ("Not Desert-Creating") and is plotted for comparison. The outcome shown here is monthly prescription drug days suply (adjusted for multi-month supplies as described in the main text). Period 0 is the period of closure. To improve readability, the plot shows the coefficients staggered slightly as a visual aid.

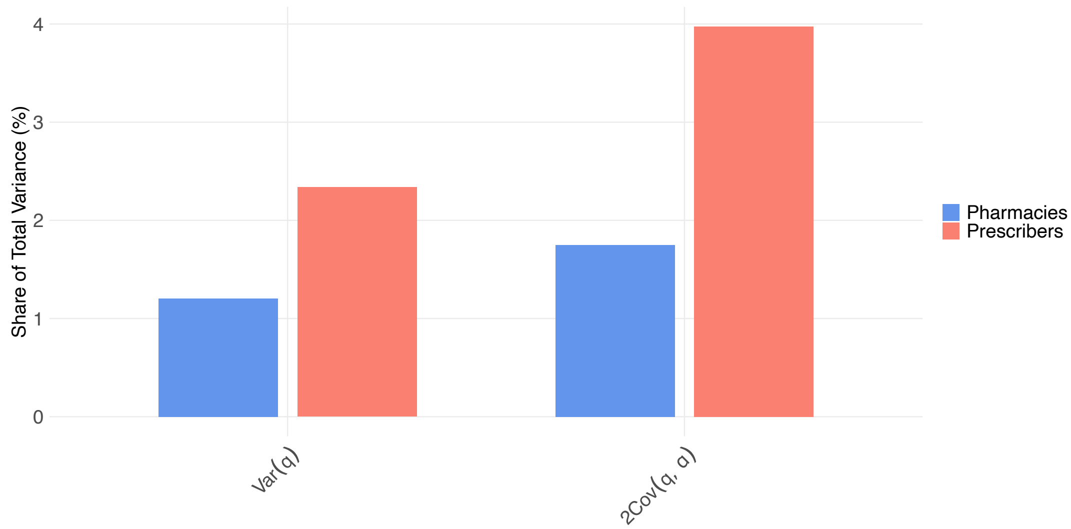

(a) Share of Overall Variance Explained by Pharmacy/Prescriber Effect and Sorting

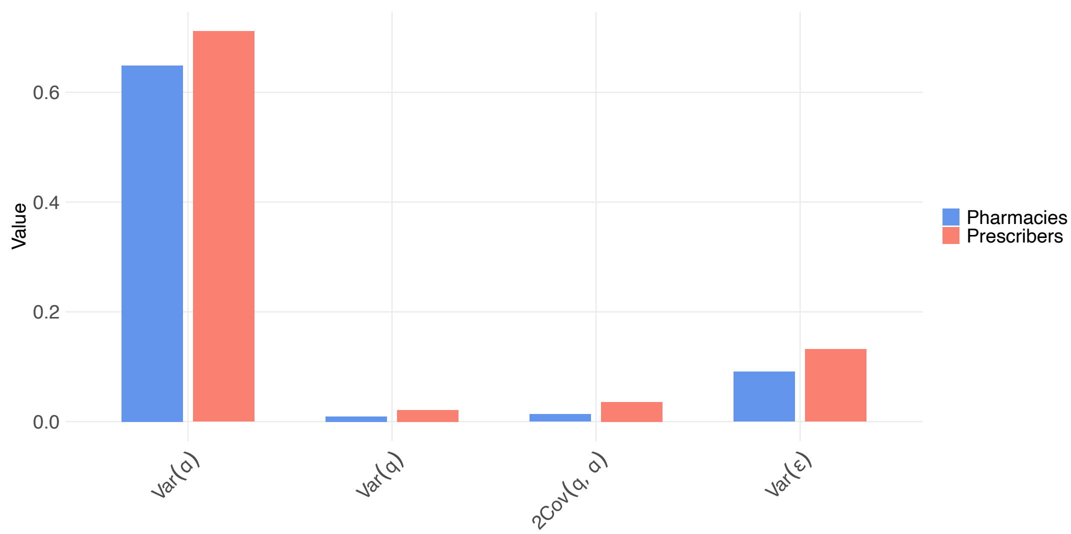

(b) Variance and Covariances from Pharmacy and Prescriber AKM Models

Figure 7: Estimates of variances and covariances of fixed effects in Equations 10 (pharmacies, in blue) and 11 (prescribers, in red). All estimates are adjusted to account for limited mobility bias using the heteroskedastic-robust correction from Kline, Saggio and Sølvsten (2020). The blue bars in Panel (a) show the variance of the pharmacy effect and the covariance of this effect with the patient-by-prescriber-by-insurer effect as a share of the total variance of that model. Similarly, the red bars in Panel (a) show the variance of the prescriber effect and the covariance of this effect with the patient-by-pharmacy-by-insurer effect both as a share of the overall variance of that separate model. Panel (b) reports the variances of each component of each model.

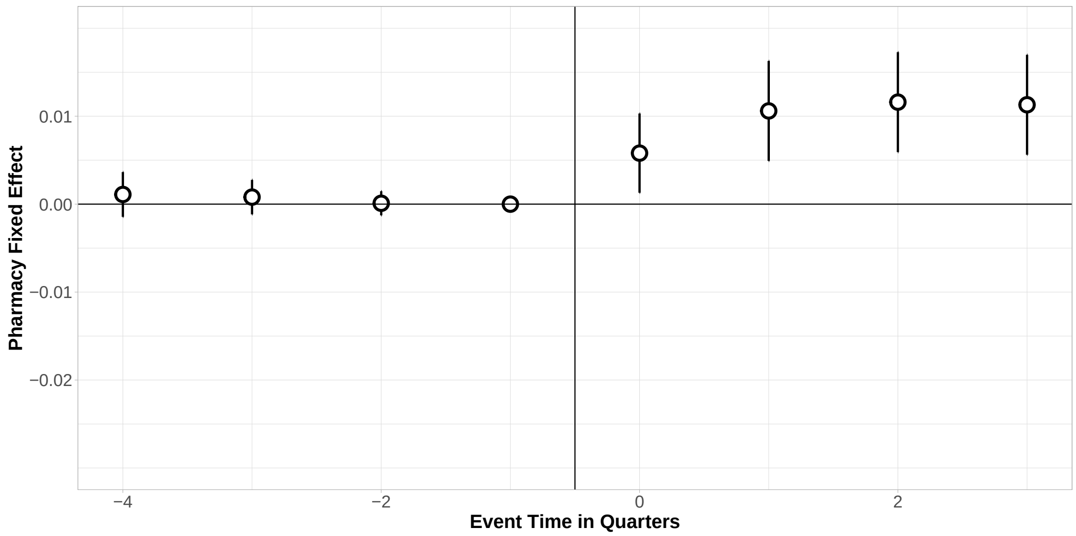

(a) Effect of Closure on Low-Income Patients' Dispensing-Weighted Pharmacy Fixed Effect

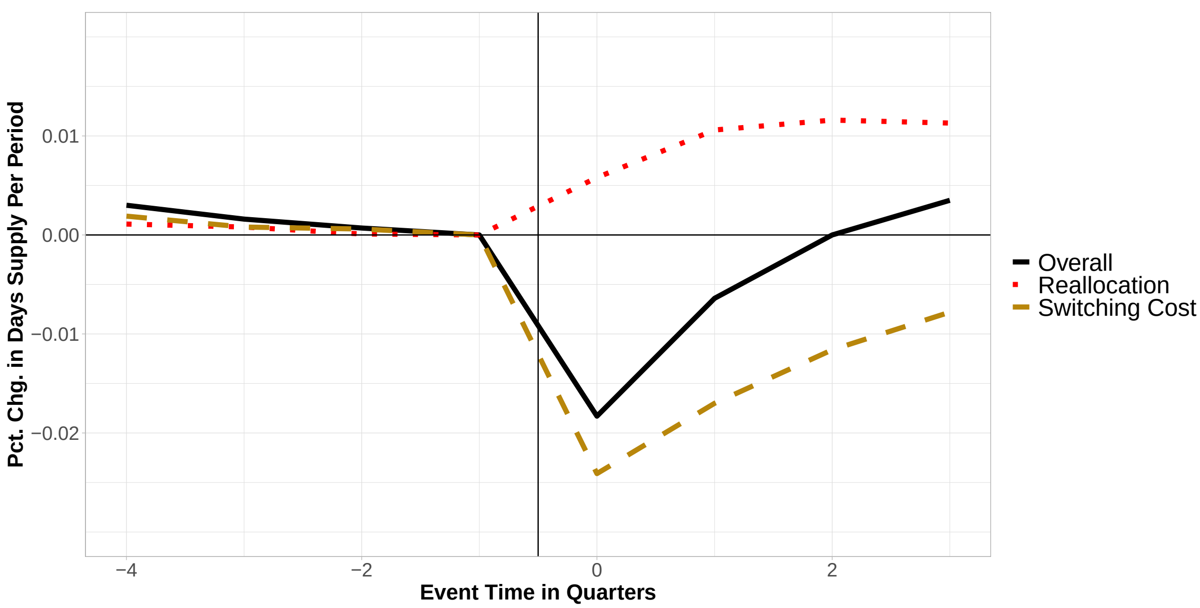

(b) Decomposing Mechanisms Driving Overall Effect of Exits for Low-Income Patients

Figure 8: Panel (a) shows the effect of closures on the model-estimated fixed effects of pharmacies that low-income patients use. This regression subsets to the set of closing pharmacies with an estimated fixed effect and then implements the standard reduced-form event study using patients experiencing closure one year in the future as controls. The fixed effects are weighted by their share of the patients' consumption that quarter. Note that the panel is unbalanced because patient-quarters with no drug consumption are dropped. Panel (b) leverages the fact that prices are low and do not change significantly to estimate the switching costs. The dotted red line just reprints the point estimates from Panel (a), while the black line is the reduced-form effect of closures for the subset with estimated fixed effects (using a Poisson specification). The brown dashed line is the switching cost term and is simply the difference between the solid black and dotted red lines.

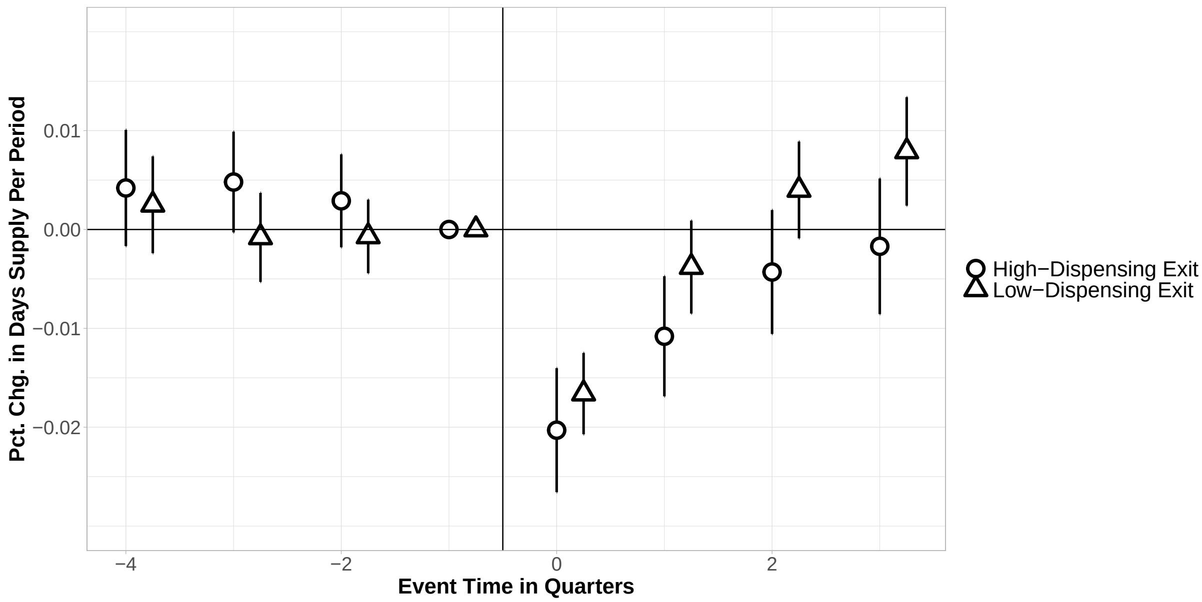

(a) Effect of Low vs. High-Dispensing Exits on Low-Income Patients' Drug Consumption

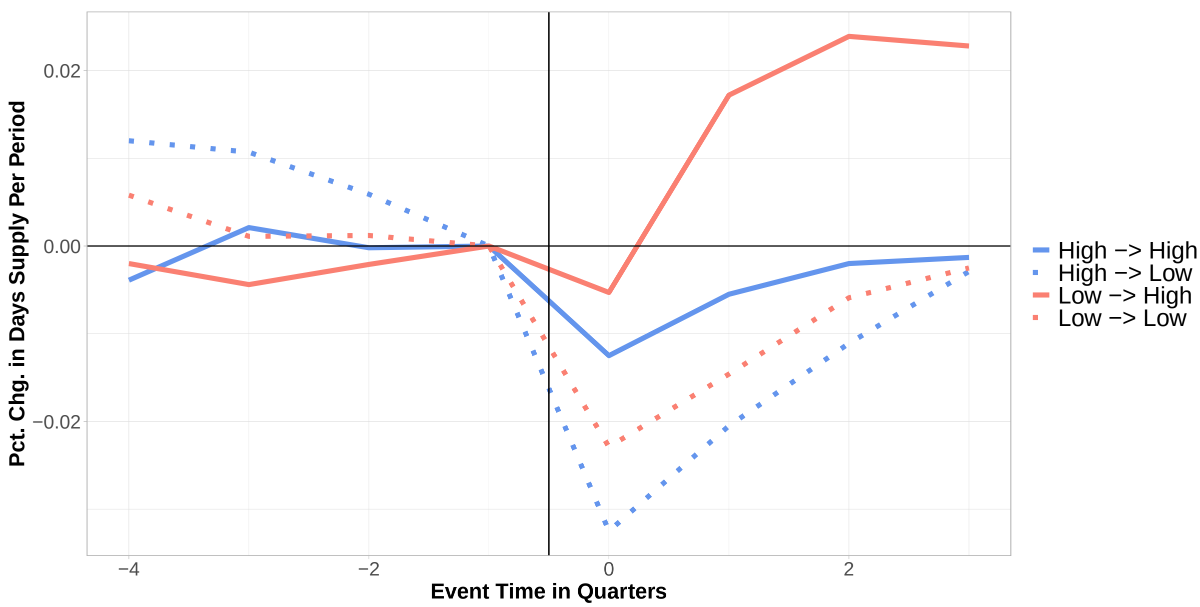

(b) Effect of Closures on Low-Income Patients' Drug Consumption by FE of Closing and New Pharmacy

Figure 9: Panel (a) shows the effect of closures split by whether the closing pharmacy is above or below the patient-weighted median pharmacy fixed effect. The plotted event study coefficients are the result of two separate implementations of Equation 8, where treated patients shopping at low-dispensing exiting pharmacies are matched to low-dispensing controls, and similarly for treated patients shoppign at high-dispensing exiting pharmacies. Panel (b) splits each of the two groups in Panel (a) into two further groups based on whether the average fixed effect of the patient's *post-period* pharmacies is above or below the median. However, because the control group does not experience an exit during the study period, I can only match treated patients to controls based on whether the pre-period pharmacy is above or below the median. This conditions on an outcome and requires dropping patients who drop out on the extensive margin (because I cannot then match them to a pharmacy fixed effect post-closure).

A Appendix Tables and Figures

	pharmacies	pharmacies	prescribers	prescribers
	full	loo-connected	full	loo-connected
Observations	18434702		18470762	
Patient-X-Insurers	10255098		8625828	
Providers	79470		841745	
Average Log Days Supply	7.11		7.11	
Non-Singleton Obs	12197597	11862876	14072259	12407762
Patient-X-Insurers	4017993	3937977	4227325	3852192
Providers	75089	66631	711011	426496
Average Log Days Supply	7.29	7.25	7.2	7.15
Patient-X-Insurers by Provider Spells	4899672	4757968	6774334	5880142
Moves / Patient-X-Insurers	0.219	0.208	0.603	0.526

Table A1: Summary statistics for the pharmacies and prescriber AKM models. The "full" columns refer to the full set of data after implementing the sample restrictions described in the text. The "loo-connected" columns describe the remaining observations after subsetting to the leave-one-out connected set. The rows labeled Patient-X-Insurers refer to Patient-Prescriber-Insurers for the first two columns and Patient-Pharmacy-Insurers for the last two. Similarly, the row labeled "Provider" refers to the number of pharmacies in the first two columns and the number of prescribers in the last two. "'Non-singleton observations" refers to Patient-X-Insurers that are in the data for more than one year. The average log days supply is at the calendar year level.

variable	Pharmacy FE
Claims (x1,000,000)	-0.017
	(0.053)
Mail Order Proxy	0.102
	(0.014)
Chain	-0.073
	(0.002)
Exit	-0.016
	(0.003)
Desert-Creating Exit	-0.017
	(0.009)
Constant	0.3
	(0.002)
N	66631
R-sq	0.022
adj. R-sq	0.022

Table A2: Estimates from regression of model-estimated pharmacy fixed effects on various pharmacy characteristics. The regression is run at the pharmacy level and is unweighted by patients or dispensing. Note that these FEs have not been shrunk to account for limited mobility bias. See Appendix B for details on how the pharmacy characteristics are defined.

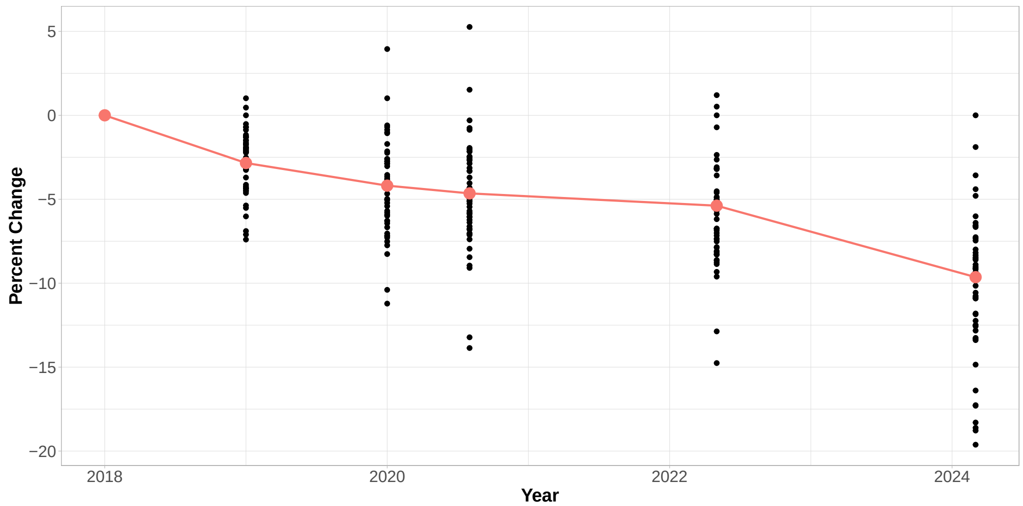

Figure A1: Change in net number of pharmacies 2018-2024. Because my claims data ends in 2019, I use public data on Part D insurers' pharmacy networks to examine changes in the number of pharmacies in more recent years. This data does not allow me to precisely determine entry and exit dates. Instead, at a point in time, a pharmacy is defined to be open if it is included in the insurance network of at least one Part D plan. Each dot represents the change since 2018 in the net number of pharmacies in a given state. The solid red line shows the average for the United States.

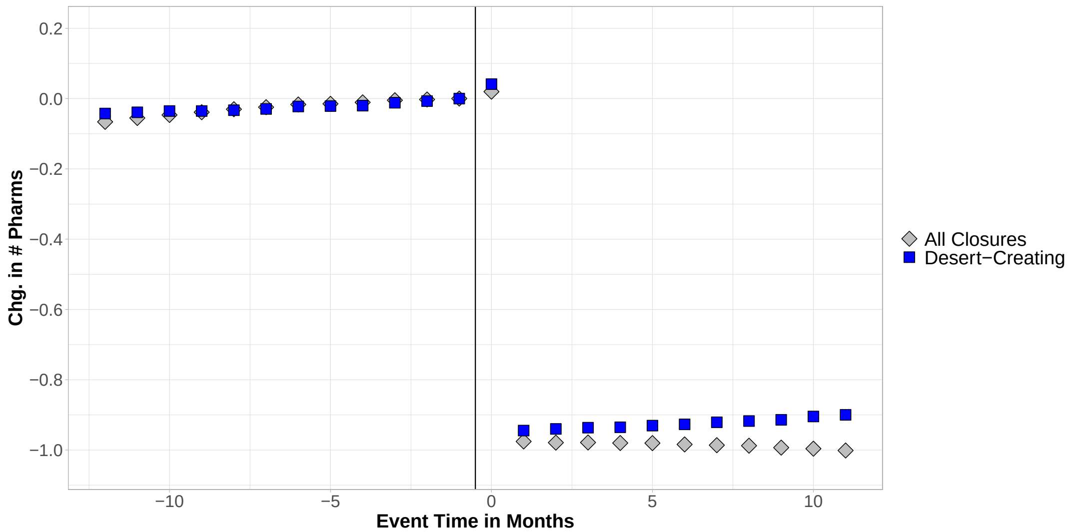

Figure A2: Net change in the number of pharmacies in the five digit Zip code of an exiting pharmacy in the two years surrounding closure. The gray diamonds show the results for all closures, while the blue squares subset to desert-creating pharmacy closures.

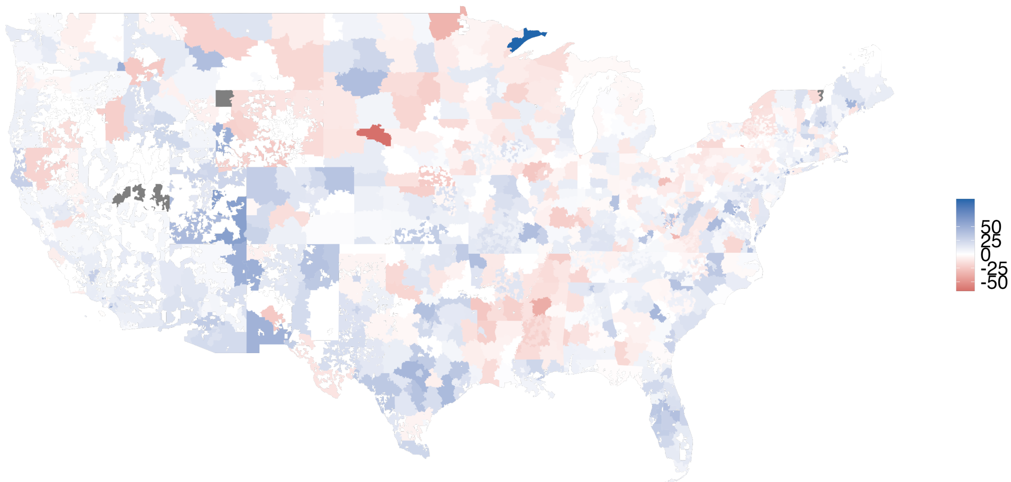

Figure A3: Change in the net number of pharmacies between 2009 and 2019 at the three-digit Zip code level. The shading corresponds to percent changes in the number of pharmacies. Missing Zip codea are gray, while Zip codes with no change are white..

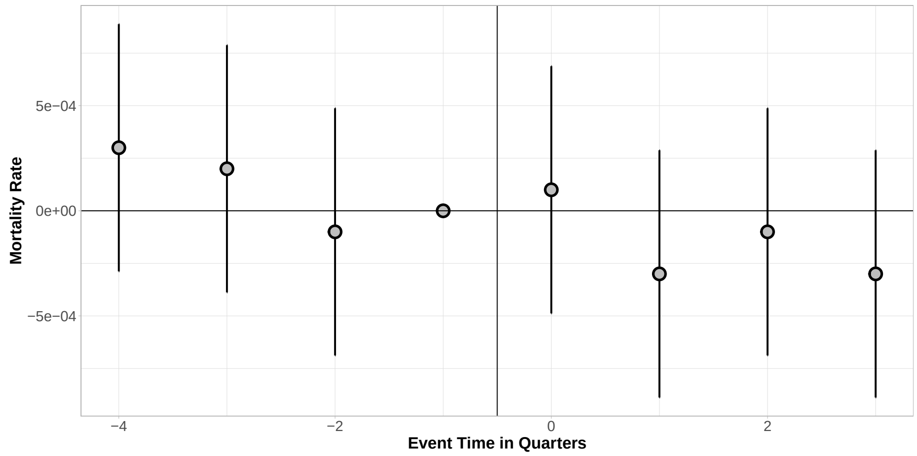

(a) Mortality

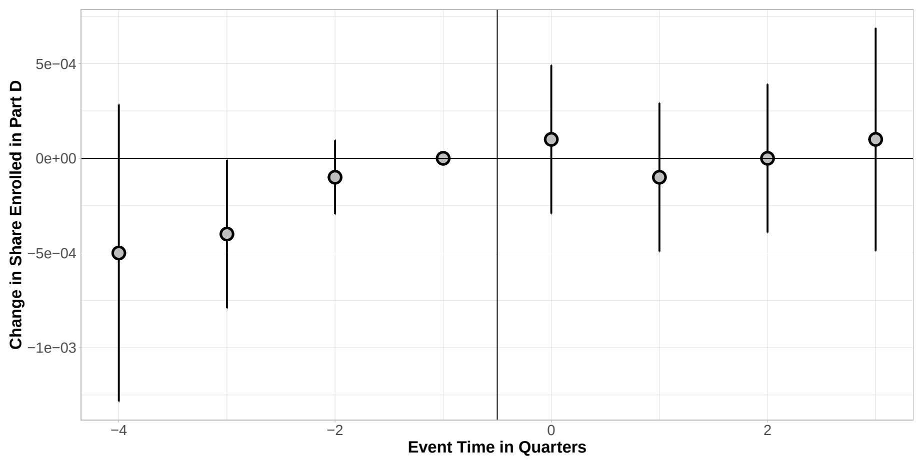

(b) Part D Enrollment

Figure A4: Event study coefficients and 95% confidence intervals from Equation 8. Unlike all other event studies, these do not condition on patients being alive and enrolled in Part D for the entire 24-month study period. Instead, Panel (a) shows the effect on mortality. Panel (b) conditions on the set of patients who are alive for the 24-month period and shows the effect on Part D enrollment. Month 0 is the month of closure.

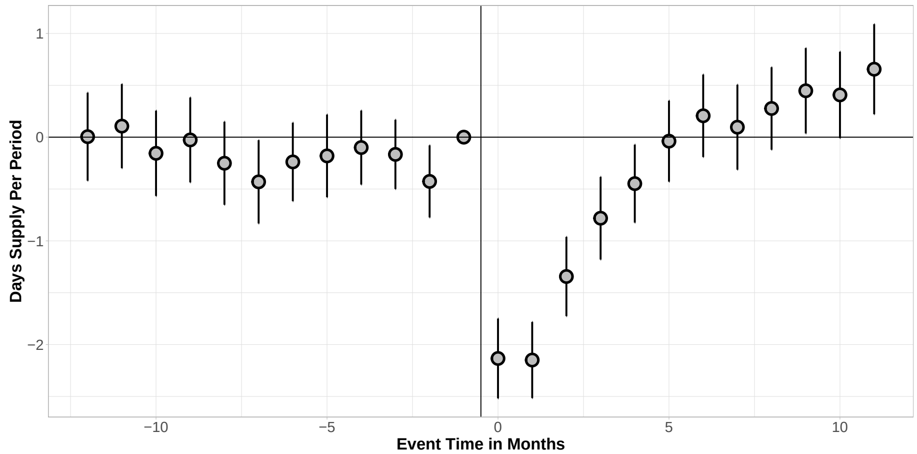

(a) All Patients, Linear Regression

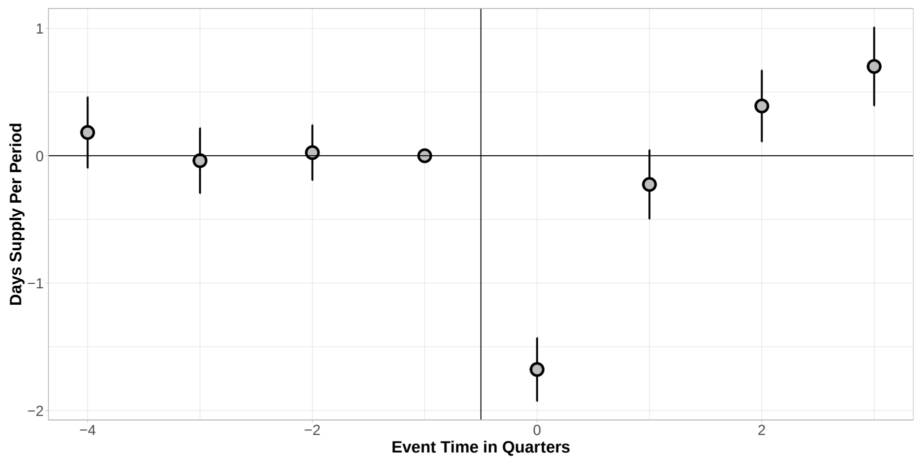

(b) All Patients, Linear Regression, Months Binned into Quarters

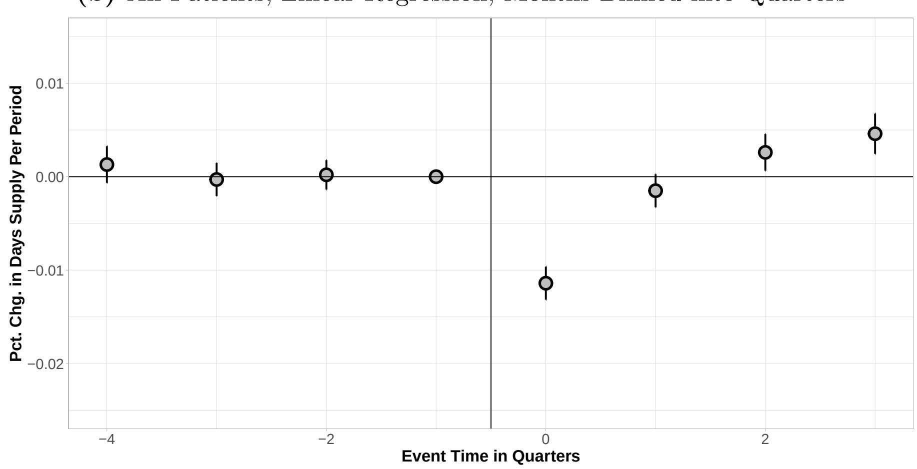

(c) All Patients, Poisson Regression, Months Binned into Quarters

Figure A5: Event study coefficients and 95% confidence intervals from Equation 8. The outcome shown here is monthly prescription drug days supply (adjusted for multi-month supplies as described in the main text). Panels (a) and (b) show the results of a linear specification, while Panel (c) shows the estimates from a Poisson regression. Furthermore, in Panels (b) and (c), months are binned into event-time quarters to increase precision but all outcomes are still kept at the monthly level). Month 0 is the month of closure.

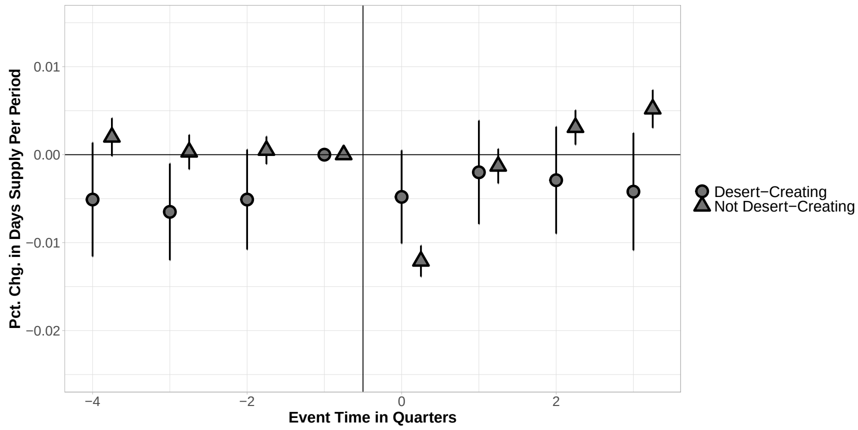

(a) All Desert-Creating Closures (Not Just Rural)

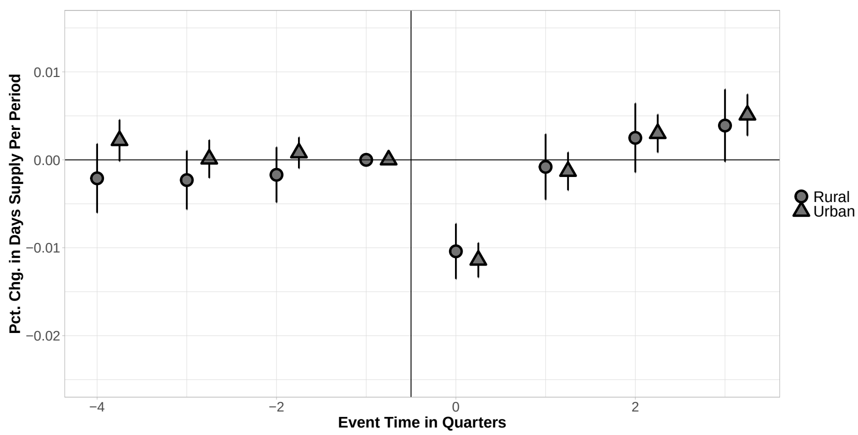

(b) Rural vs. Urban Pharmacy Closures

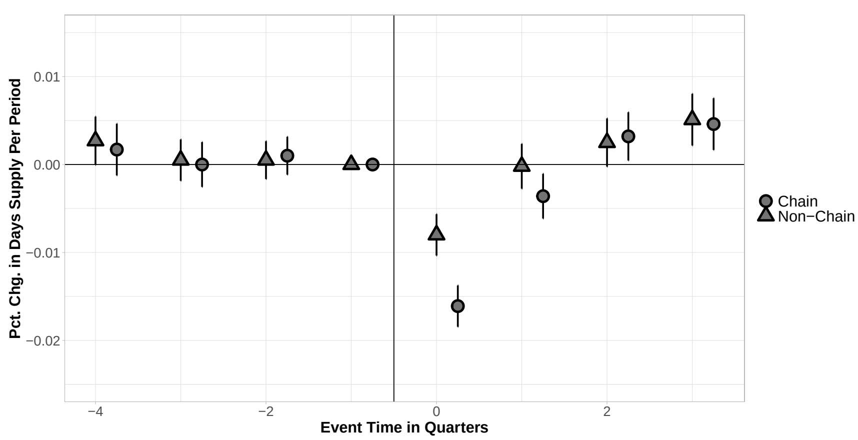

(c) Chain vs. Independent Pharmacy Closures

Figure A6: Event study coefficients and 95% confidence intervals from Equation 8 using the Poisson specification and splitting patients based on characteristics of their closing pharmacy (as defined in the text and in Appendix B). The outcome shown here is monthly prescription drug days suply (adjusted for multi-month supplies as described in the main text). Period 0 is the period of closure.

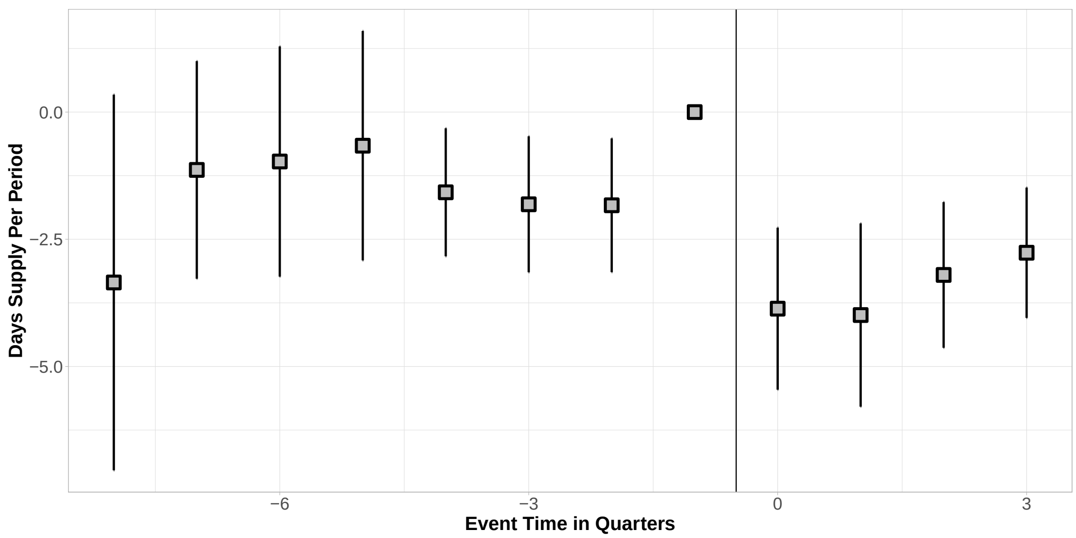

(a) Chain Pharmacy Network Exclusions, Same Insurer Controls

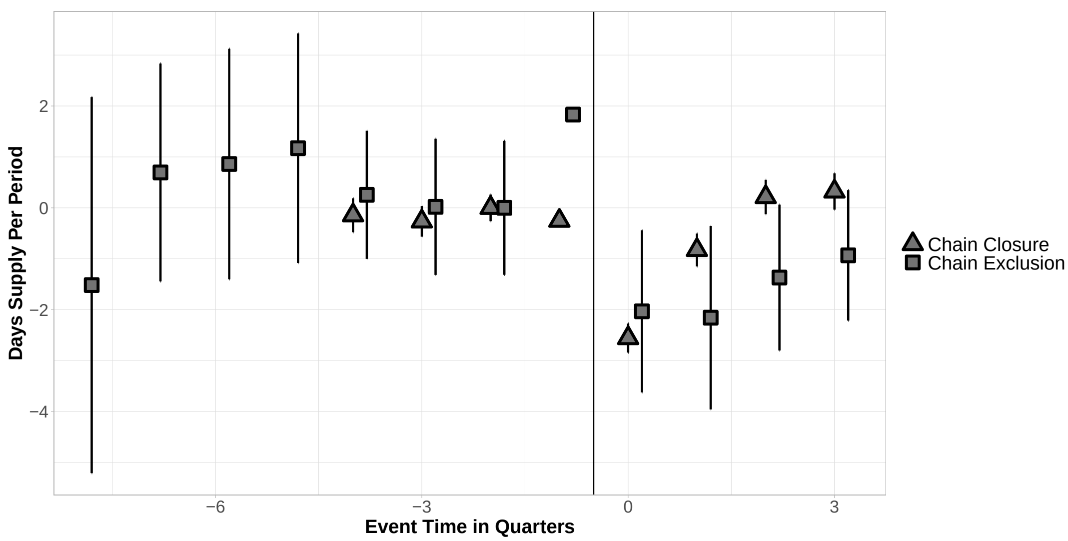

(b) Chain Pharmacy Network Exclusions Compared to Chain Closures

Figure A7: Event study coefficients and 95% confidence intervals from Equation 8. The outcome shown here is monthly Medicare Part D drug days supply. Panel (a) shows the results for the chain pharmacy network exclusion analyses described in Section 4.6. Panel (b) compares the network exclusion results to chain pharmacy exits. For interpretability, period -2 rather than -1 is normalized to an effect size of 0. Month 0 is the month of closure/exclusion. Note: Updated version of this figure using a Poisson model is forthcoming.

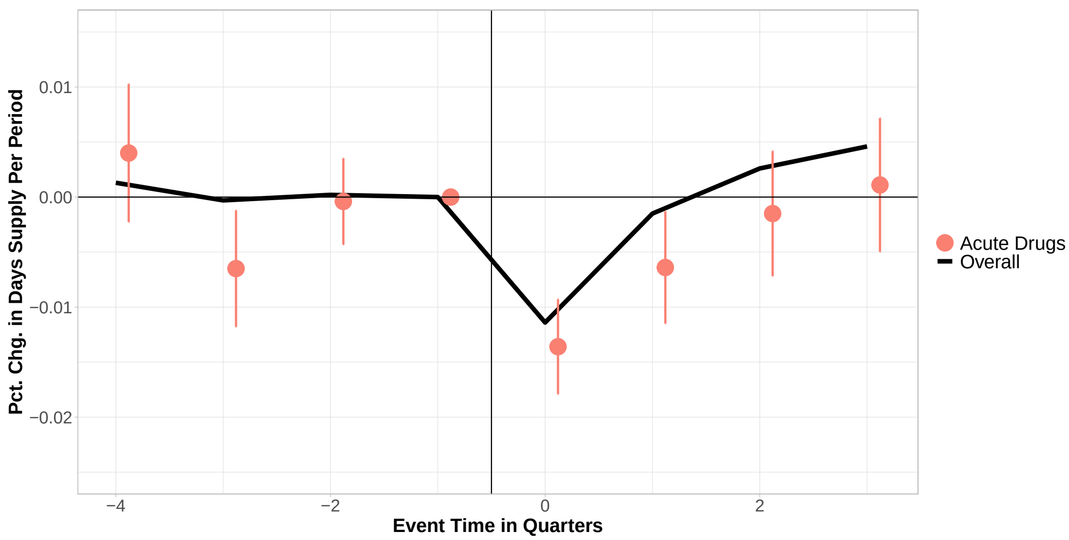

Figure A8: Event study coefficients and 95% confidence intervals from Equation 8 using the Poisson specification on the full set of patients. The outcome shown here in red is monthly prescription drug days suply of drugs that treat acute conditions, compared to the overall (the black line). Period 0 is the period of closure.

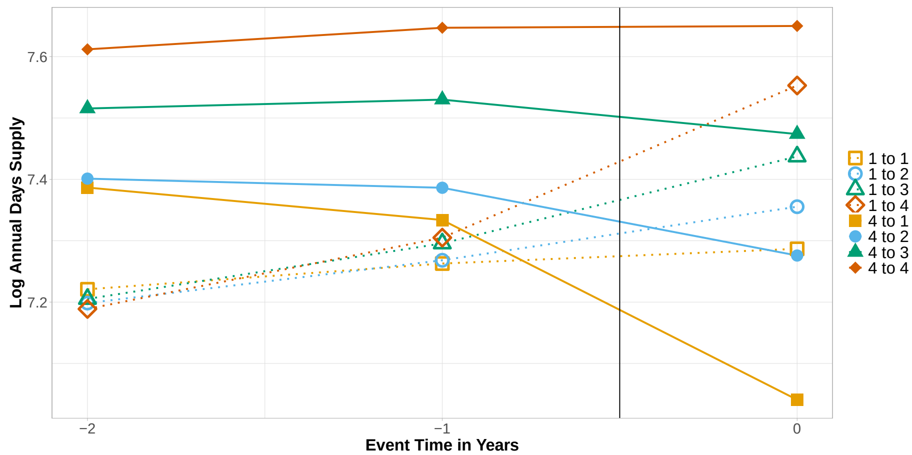

(a) Symmetry Test: Pharmacy Regression

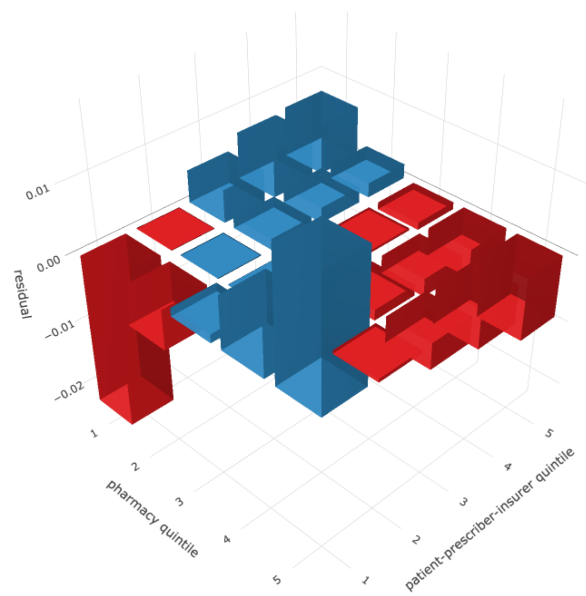

(b) Residual Test: Pharmacy Regression

Figure A9: These plots show two validation checks for the AKM regression to recover the pharmacy fixed effects. Panel (a) subsets to the set of moving patients with two years of pre-period data and who start at a pharmacy with a fixed effect in either the bottom (1) or top (4) quartile. These two groups of patients are further subdivided into four groups based on the quartile of the pharmacy they move to in year 0. The vertical axis shows log annual drug days supply. Remember the patient is actually a patient-prescriber-insurer combination. Panel (b) shows average residuals of the AKM model split into 25 groups based on quintiles of the pharmacy and patient-prescriber-insurer fixed effect.

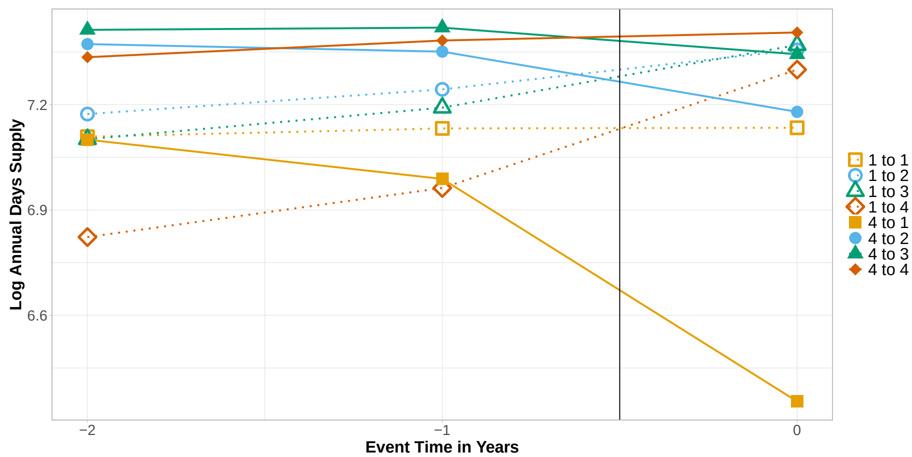

(a) Symmetry Test: Prescriber Regression

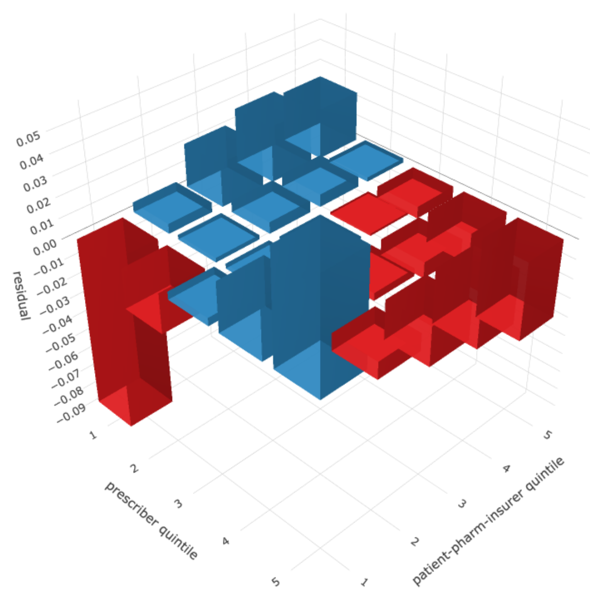

(b) Residual Test: Prescriber Regression

Figure A10: These plots show two validation checks for the AKM regression to recover the prescriber fixed effects. Panel (a) subsets to the set of moving patients with two years of pre-period data and who start matched to a prescriber with a fixed effect in either the bottom (1) or top (4) quartile. These two groups of patients are further subdivided into four groups based on the quartile of the prescriber they move to in year 0. The vertical axis shows log annual drug days supply. Remember the patient is actually a patient-pharmacy-insurer combination. Panel (b) shows average residuals of the AKM model split into 25 groups based on quintiles of the prescriber and patient-pharmacy-insurer fixed effect.

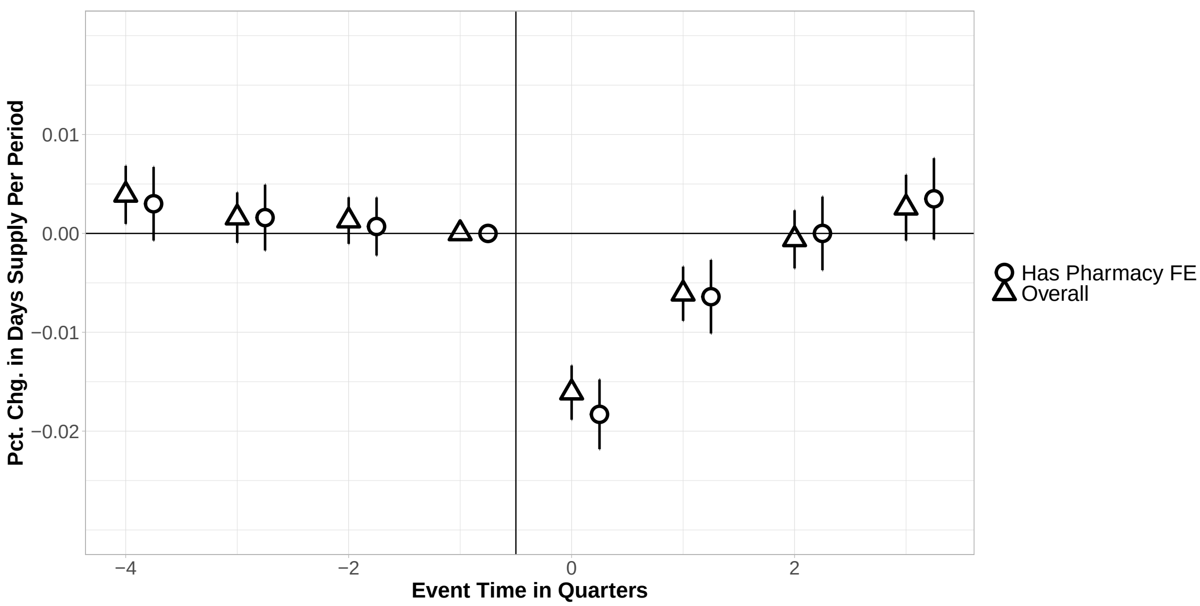

Figure A11: Event study coefficients and 95% confidence intervals from Equation 8 using the Poisson specification on the set of low-income patients. The figure shows the results of two separate regressions, one run on patients for whom the closing pharmacy's fixed effect is defined by the AKM model and the other on the overall set of low-income patients experiencing a closure. The outcome is monthly prescription drug days suply (adjusted for multi-month supplies as described in the main text). Period 0 is the period of closure.

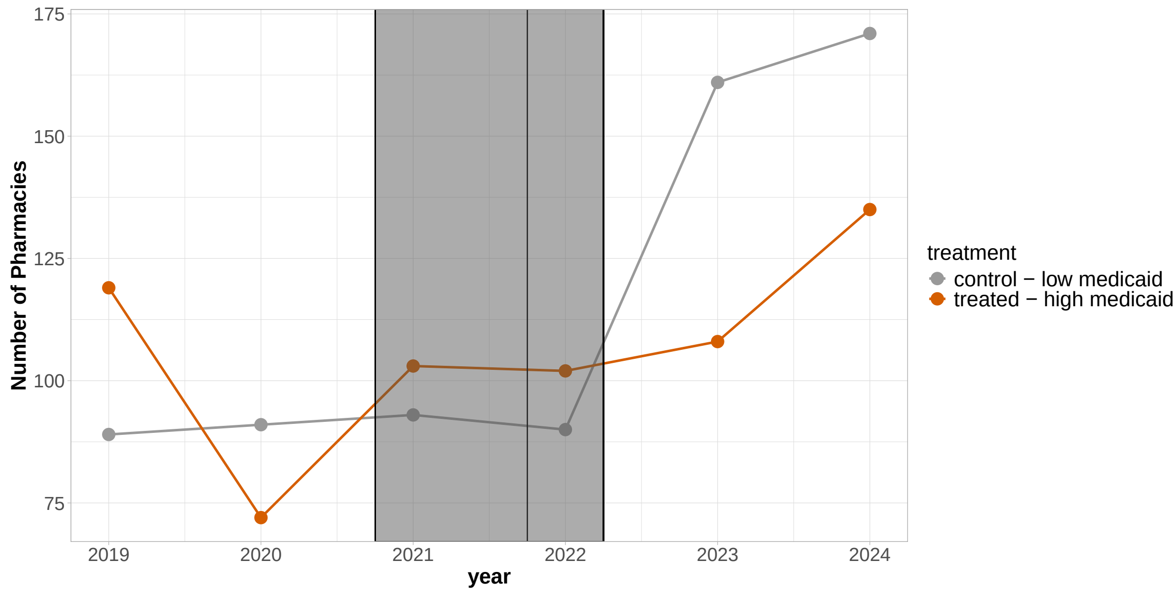

(a) CA Pharmacy Exits

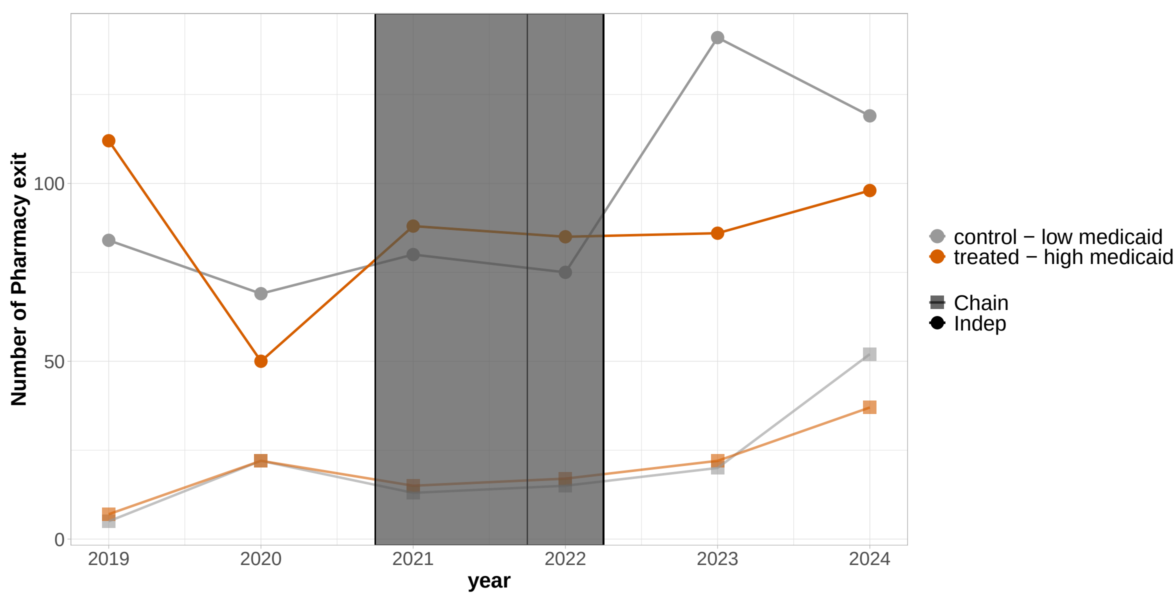

(b) CA Pharmacy Exits, Chain vs. Indep.

Figure A12: Raw counts of annual pharmacy exits in California. The shaded region represents the fact that the state planned to carve out prescription drugs from its Medicaid managed care contracts in 2021 before delaying the policy to 2022. Treated pharmacies are defined as those in above-median Zip codes in terms of the Medicaid population share and are indicated in orange, while controls are in below-median Zips and are denoted in gray. Panel (a) shows the aggregates, while Panel (b) additionally splits out exits of chain pharmacies (square markers) versus independent pharmacies (circular markers).

B Harmonizing Pharmacy IDs and Defining/Categorizing Exits

The "treatment" of interest in the reduced form analyses is physical exits of pharmacy locations. To identify these exits, I start with the identifiers included as the "dispenser" on Part D claims. From 2009 onwards, these identifiers are the dispensers' national provider identifier (NPI) about 99% of the time. In 2008 and especially 2007 a significant share of claims were filed using a different dispenser identifier (the NCPDP ID). To avoid conflating switching from the NCPDP ID to the NPI as an exit, only NPIs are used to define exits.

NPIs are flagged either as individuals or institutions. Pharmacies are institutions, but in some cases, an individual doctor may dispense a drug, in which case they would be recorded as the dispenser. For the purposes of defining entries and exits, I only consider institutional NPIs.

Because a given pharmacy location can have multiple NPIs that can change over time, I merge the NPIs to pharmacy location data from the National Plan and Provider Enumeration System (NPPES). This location data includes addresses, which I clean and standardize and then combine all NPIs dispensing from that address in the same or adjacent months into a single "harmonized NPI". This harmonized NPI thus captures pharmacies with multiple NPIs or cases where an NPI changed, but not the pharmacy location (e.g., following a change in ownership).

I then define entry and exit as the first and last date one of these harmonized NPIs filled a Part D claim. To avoid conflating the effects of entry and exit, I only consider exits where the harmonized NPI opened at least 24 months prior. This restriction also excludes some very small dispensers that briefly pop up in the data for short periods of time as counting as "exits". In all of these cases, these exclusions restrict the set of pharmacies that I define as exits, but are all kept in the data when defining patients' outcomes.

Mail-Order: While patients do not travel when they get their drugs via mail-order, the NPIs are still defined based on the location of the facility where the prescriptions are filled. These exits may have very different effects on patients, especially because the notion of a "closure" is unclear in these cases (e.g., a company choosing to move prescription fills to another facility may be recorded as a "closure" but from the patients' perspective, nothing changed). This fact combined with the policy focus on retail exits, lead me to exclude considering mail-order exits.[57] However, in early years of the data, I lack flags for mail-order pharmacies. Thus, I use the fact that mail-order pharmacies tend to be large to leverage a dispensing threshold to drop mail-order pharmacies. I err on the conservative side and drop only pharmacies that fill more than 50,000 claims in any year prior to exit. (Because I have a 20% sample, this means that the pharmacy filled 250,000 Part D claims alone.) I validate this definition in later years of data when I have a flag for mail-order pharmacies and confirm that this definition means I miss some mail-order pharmacies, but misclassify very few retail pharmacies as mail-order.

Rural vs. Urban: Pharmacies are defined as urban or rural using the definition used by the Federal Office of Rural Health Policy. By this definition, rural areas are those outside of metropolitan areas, which correspond to Rural-Urban Commuting Area Codes of 4 through 10. I assign these codes based on a pharmacy's Zip code.

Chain vs. Non-Chain: I start by defining chain pharmacies using flags provided in the 2013 file of data from the National Council on Prescription Drug Providers (NCPDP). However, because these flags are incomplete, I supplement with additional hand-collected chain names by examining the NCPDP and NPPES for frequently repeated business names and validating them as chains. I then define the set of chains in the NPPES as the set of pharmacies where either any of the business names match one of these chain names.

"Desert-Creating" Closures: I define desert-creating closures as those where the closing pharmacy dispensed at least half of the days supply among all non-mail order pharmacies dispensing in that Zip code **and** of the patients living in that Zip code in the 7-12 months prior to exit. The first restriction ensures that there are not other substantial sources of supply in the same Zip code. The latter helps eliminate mail-order pharmacies. I use the Zip code as the geographic unit over which deserts are defined to match the finest geographic data Medicare provides on patients. While arbitrary, I consider the 50% threshold generous in that it classifies as "desert-creating" cases where there may be other (albeit smaller) pharmacies nearby. However, these types of closures are still rare, accounting for less than 11% of the patient-weighted total.

[57] Of course, as stated above, these observations are kept in the data when measuring patients' outcomes. Indeed, switching to mail-order is one potential response to experiencing a retail exit.

C Details on Drug Classification

The Part D data provide the national drug code (NDC) and brand name for each drug claim. By merging to the formulary files, I also obtain the generic name. I use these identifiers to merge to the Micromedex RedBook classification, which provides a therapy class and group for each drug, as well as an indicator for whether the drug is primarily used to treat chronic conditions, acute conditions, or both. I only have access to this classification system through the end of 2013, but by merging to the generic names, I am able to extend the classification through the end of my study period. The only drugs I am unable to classify are new drugs introduced after 2013, but these account for a small share of dispensing. Furthermore, because the controls are drawn from the same periods as the treateds in the closure regressions, any increasing non-match rates should affect both groups similarly and thus not impact the treatment estimates.

Chandra, Gruber and McKnight (2010) Classification: The authors enlist a panel of doctors and pharmacists to classify drugs based on how quickly discontinuation can lead to an adverse event (e.g., hospitalization). The three categories are: drugs where discontinuation is likely to cause an adverse event within a month, within a year, or greater than a year. I use code from Lavetti and Simon (2018) to connect the Chandra, Gruber and McKnight (2010) groups to Redbook.

"High-Value" Cardiovascular Drugs: I follow Choudhry et al. (2011) in considering four classes of cardiac drugs: statins, beta blockers, ACE Inhibitors, and Angiotensin-Receptor Blockers (ARBs). These drugs treat either high blood pressure or high cholesterol and reduce the risk of adverse events like heart attacks and strokes. Because these conditions are chronic, these drugs are typically taken long-term. Most drugs in these classes are now available as cheap generics and are among the most consumed drugs in Part D.

Acute Drugs: Using the Redbook classification, I subset to acute drugs in Figure A8 as a robustness check to address stockpiling and intertemporal substitution concerns. Because acute drugs are taken for short periods of time, patients are less likely to have a stock of these drugs. Thus, the fact that they show a similar pattern in use as other drugs after closure provides evidence against the argument that the declines following closure are due to patients consuming a stock of pills.

D Policy-Driven Changes in Pharmacy Reimbursements

Despite increasing regulations on insured patients' drug copays and coverage levels since the roll out of Medicare Part D, the government's role in setting pharmacy reimbursements had been declining until very recently. The roll-out of privatized prescription drug benefits in Medicare via Part D in 2006 and the privatization of many state Medicaid benefits in capitated managed care programs both increased the role of insurers and reduced the role of the government agencies in negotiating reimbursements with pharmacies. However, since 2019, this trend has reversed for prescription drugs, with six states "carving out" prescription drugs from their Medicaid managed care contracts (Bendicksen and Kesselheim, 2022). These carve-outs mean that the state typically sets Medicaid pharmacy reimbursement rates directly, rather than allowing a managed care insurer to negotiate reimbursements with pharmacies. While administratively-set reimbursements need not necessarily be higher than negotiated rates, prior work has shown that this is overwhelmingly the case (Dranove, Ody and Starc, 2021). More directly, a small but growing number of states have recently imposed reimbursement floors that pharmacy benefit managers must pay pharmacies in their states.[58]

I present a descriptive case study of one state Medicaid drug carve out with specific features that help isolate the effect of reimbursements on pharmacy exit. California carved out prescription drugs from its state Medicaid managed care program in 2022. Thus, rather than negotiating with insurers, pharmacies started receiving the reimbursements set by California Medicaid. These administratively set reimbursements included a \$3.15 supplement to the dispensing fees paid to small and medium pharmacies with less than $90,000$ total annual fills. Under the assumptions that private insurers were not providing higher reimbursements to smaller pharmacies,[59] and parallel trends between small and large pharmacies, comparing pharmacies below the threshold to those above allows differencing out non-reimbursement changes arising from the carve-out.

I thus show in Figure A12 four groups of pharmacies split by whether they were in zip codes with above or below the state median share of Medicaid recipients and by whether the pharmacy is a chain or non-chain. I use the chain/non-chain split as a proxy for pharmacy size, because independent pharmacy locations dispense about 35% as many fills as chains on average Ladsariya

[58] States that recently imposed these laws include Kentucky, Tennessee, West Virginia, and Alabama.

[59] This is unlikely to be the case. Indeed, much of the recent policy focus has centered on claims that small independent pharmacies get *lower* reimbursements.

et al. (2023). Because all of the policies considered were implemented after the claims data end in 2019, pharmacy openings and closings are defined using CMS public-use files, as described in Section 2.3.

Figure A12 shows the number of annual pharmacy exits in the six years surrounding the policy change in California. In the years preceding the policy change, the rates of closure across high and low-Medicaid pharmacies were broadly similar. However, in the two years following the policy change, a gap emerges, with the control group (i.e., pharmacies in low-Medicaid zip codes and thus less exposed to the policy change) experiencing a 20% higher rate of closure relative to treated (high Medicaid zip code) pharmacies. Panel (b) of Figure A12 shows that this gap is especially large among independent pharmacies, which tend to be small, and thus likely to receive the higher reimbursement rates that California provided to smaller pharmacies.

Separating the costs of the reimbursement change from the overall carve-out is challenging. Thus, to estimate the costs of this policy, I focus on the difference between the chain and independent pharmacies. I make four strong assumptions. First, I assume that only the independent pharmacies are getting the $3.15 small-pharmacy supplement. Second, that these pharmacies were getting the same reimbursements as chains in the pre-period when reimbursements were negotiated with insurers. Third, that any other changes associated with the carve out impacted these two groups evenly. And finally, that the state paid the supplement on 10-30% of the 145 million Medicaid prescription drug fills in 2022-2023.[60] With these assumptions, the difference in the rates of closure between the two groups (30 pharmacies per year in 2023 and 2024) is caused by the $3.15 supplement, and comes out to roughly $1.5-4.6 million per year per averted closure.

[60] https://hcai.ca.gov/visualizations/healthcare-payments-data-hpd-snapshot/.