Causes of Retail Pharmacy Closures

Chanwool Kim*

This draft: July 30, 2026
Draft updated frequently: please click here for the latest version.

Job Market Paper

Preliminary and incomplete. Please do not circulate.

Abstract

Recent retail pharmacy closures have raised concerns about "pharmacy deserts" and access to prescription drugs, prompting state and federal proposals to raise the margins pharmacies earn on dispensing. This paper shows that the closures stem less from declining drug profitability than from the collapse of the front-store retail business that historically cross-subsidized dispensing. Using a novel dataset linking pharmacy sales, drug prices, and consumer shopping behavior, I document that the U.S. pharmacy market comprises two business models: multi-product chain pharmacies ("chains") that rely on consumer packaged goods (CPG) sales to offset low prescription (Rx) margins, and small independent pharmacies ("independents") that focus almost exclusively on dispensing. I demonstrate four empirical facts: (i) exits are concentrated among chains; (ii) exits reduce displaced patients' drug adherence only temporarily; (iii) exits concentrate where pharmacies are dense, leaving distance and driving time to the nearest pharmacy essentially unchanged; and (iv) exit rates track declining CPG sales, not falling Rx margins. Because entry and exit are forward-looking decisions that respond to total store profitability rather than the drug margin alone, I estimate a multi-category demand model that converts dispensing and front-store channels into store-level profit dollars and embeds it in a dynamic model of entry and exit. The model projects that the number of pharmacies will shrink by nearly half in the long run under current market conditions. Still, because projected closures are concentrated in dense markets, the distance to the nearest pharmacy barely changes. The consumer cost is modest and borne mainly through the prices charged by surviving stores rather than by lost access. My estimates imply that restoring margins to their real-terms peak would cost $30 billion per year, while recovering only a tenth of the decline and returning about one cent of consumer surplus per public dollar.

*Kenneth C. Griffin Department of Economics, University of Chicago, 1126 East 59th Street, Chicago, Illinois 60637, United States (email: chanwoolkim@uchicago.edu). I am grateful for my advisors, Jean-Pierre Dubé, Ali Hortaçsu, and Milena Almagro, for guidance and encouragement. I also thank Dan DiGiorno, Joseph Dunham, Arjun Gopinath, Joshua Gottlieb, Shanon Hsuan-Ming Hsu, Youngeun Lee, Casey Mulligan, Ricardo Quineche, Estéfano Rubio, Francesco Ruggieri, Camilla Schneier, Kyle Shinn, Younggeun Yoo, and participants at the University of Chicago Microeconomics Third Year Seminar, Public and Labor Seminar, and Industrial Organization Seminar for helpful comments and discussion. Jooyoung Lee provided outstanding research assistance. Researcher's own analyses calculated (or derived) based in part on data from Market Track, LLC dba Numerator and marketing databases provided through the Numerator Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Numerator data are those of the researcher and do not reflect the views of Numerator. Numerator is not responsible for and had no role in analyzing or preparing the results reported herein. Merative supplied certain data as part of one or more MarketScan Research Databases. Any analysis, interpretation, or conclusion based on these data is solely that of the author and not Merative. I gratefully acknowledge financial support from the University of Chicago Kenneth C. Griffin Applied Economic Incubator.

1 Introduction

Over the past decade, the number of retail pharmacies in the United States has fallen sharply. The contraction has raised public concern about prescription drug access and the resilience of the healthcare delivery system (Guadamuz et al., 2021a; Nowosielski, 2024). Policymakers have responded with congressional hearings, Federal Trade Commission (FTC) investigations, and a wave of state and federal proposals (Qato et al., 2019; Federal Trade Commission, 2024; NASHP, 2023; Feldman, 2023; Brown and McCuskey, 2020).

Nearly all of the proposed policies operate on a single lever: the rate at which payers and their pharmacy benefit managers (PBMs) reimburse pharmacies for dispensing prescriptions. The implicit hypothesis is that falling drug margins cause closures, leading to a conjecture that restoring those margins would keep pharmacies open in places where access is at risk.

The premise of this hypothesis is well documented. The National Community Pharmacists Association (NCPA)'s annual industry surveys and the FTC's report on PBMs both record a sustained decline in the reimbursement rates (NCPA, 2025; Federal Trade Commission, 2024). Patients see almost none of this decline as out-of-pocket payments per prescription have remained roughly flat, so the squeeze falls between payers and pharmacies. In this respect, the policy narrative is correct, and dispensing has become less profitable.

A second decline, however, is equally pronounced. Large chain pharmacies such as CVS, Walgreens, and Rite Aid have operated a hybrid business model that pairs low-margin prescriptions with high-margin consumer packaged goods (CPG). For these chains, prescriptions generate trips, and profits are earned on the soda, cosmetics, and batteries purchased alongside the fill. Indeed, all three major chains above have acknowledged the decline in front-store sales (CVS Health Corporation (2024), Walgreens Boots Alliance (2024), Rite Aid Corporation (2024)). Since both declines are real, the extent to which each contributes to the closures determines whether the proposed policies can work.

This paper therefore asks that question and its policy implication: do pharmacies close because dispensing became unprofitable or because the front store collapsed? Would restoring dispensing margins keep pharmacies open where access is at risk, at a cost commensurate with what consumers gain?

I first establish four empirical facts using a novel database that links pharmacy sales, drug prices, and consumer shopping behavior. Each of these facts runs counter to the ongoing hypothesis that reimbursement is driving closures and leading to the alleged growth of "pharmacy deserts," areas with few or no pharmacies that patients must travel far to reach. First, net closures are concentrated among the multi-product chains, not among the drug-focused independents that rely most heavily on dispensing income. Second, closures disrupt the patients they displace only temporarily: in a stacked event study, displaced patients' days of prescription supply dip by roughly 10 percent at the trough and return to the pre-closure path within a quarter. Additionally, market-level prescription consumption does not respond to net changes in local pharmacy availability. Third, exits concentrate in places dense with pharmacies, leaving the distributions of distance and driving time to the nearest pharmacy essentially unchanged. Fourth, net pharmacy losses track declining CPG purchases, not dispensing margins within counties over time.

Together, the facts imply that the closure wave has largely removed redundancy in already-served markets rather than eliminating local monopoly, and that its timing and geography follow the front store sales decline rather than the drug margin. They also affect the welfare question. As access to prescription drugs is the primary policy concern and is largely refuted by the data, what closures put at stake is secondary, though still important—travel, one-stop convenience, and the retail prices consumers pay at the stores that remain.

The facts reframe the problem, but the reduced-form evidence still cannot be naïvely taken to evaluate the proposed policies. The point is sharpest as a thought experiment. If reimbursement had been the only primitive that declined, with stores, retail prices, and preferences unchanged, no consumer choice would have moved, and the subsidy to this margin would have affected entry and exit decisions only through variable profit. However, the observed reallocation of CPG purchases away from the chains is substantial and cannot be explained solely by the reimbursement margin. Thus, which stores are saved by per-fill subsidies and where entry becomes viable depend on the joint distribution of fill volumes, retail revenues, and fixed costs. This relationship requires an understanding of the causes of the shift in CPG demand and its relationship to drug demand. In fact, retail competition is not only over drugs but also over CPGs, which naturally expands the chain pharmacies' competitors beyond other pharmacies to other CPG retailers such as dollar stores and wholesale clubs as well. Because copays are set by plan design and have remained flat, changes in reimbursement policies would affect consumers only through store choices and what they charge for everything else, and pricing those channels requires a demand system that converts changes in the choice set into dollars.

The demand model follows the multi-category competition (Thomassen et al., 2017). Consumers face shopping costs and choose which stores to visit and how much to buy in three retail categories—fresh groceries, packaged groceries, and other household essentials. Prescriptions enter the demand system as a binary category triggered by a latent health need, filled with certainty at a store the consumer chooses, and priced by payers rather than stores. The deterministic-fill treatment is motivated by event-study evidence indicating that fills are relocated rather than destroyed. The model separates the two mechanisms that can tie drug traffic to retail profit: a complementarity in preferences, whereby a fill raises the marginal utility of the rest of the CPG categories, and a shopping-cost channel, whereby the fill is a loss-leader that anchors the trip, and the front store purchase is driven by unwillingness to make a separate trip for CPGs.

The marginal utility of income, the cost of distance, and the drug–retail linkage capture consumer preferences that matter for welfare and profit calculations, and each faces a familiar endogeneity concern. First, retail prices depend on unobserved factors that affect demand. Since most retailers use national pricing, I maintain the assumption that prices are conditionally independent of unobserved taste shocks after including retailer-firm-category fixed effects to control for market-level unobserved quality and quarter dummies to control for market-level quarter effects. Second, distance enters shopping costs, but I exclude it from variable utility to help identify the parameters in shopping costs separately from those in variable utility. Third, cross-category correlations in purchases may reflect correlated tastes rather than complementarity, in which a consumer who likes shopping buys more of everything. To address this, the consumer-level panel moments separate permanent tastes from trip-level shocks by comparing within- and cross-occasion covariances. The drug interactions are read off the same consumer's retail basket on fill versus non-fill occasions at the same store, with pharmacy-format intercepts absorbing payer steering of where fills land.

To test for time-varying preferences, I estimate demand parameters separately for each year from 2019 to 2024. Taste parameters remain stable over this period. If anything, they run against a taste-based bundling, implying that the store that captures the trip captures the basket. The price coefficient is stable in *nominal* dollars across a period of 23 percent cumulative inflation, which means the marginal utility of a constant dollar rose—by roughly a fifth with the 2021 inflation surge—so consumers came out of the high-inflation era measurably more sensitive to prices than they entered it. Using the demand estimates, I vary preferences, the retail landscape, and own retail prices across years, with the largest and smallest numbers of pharmacies, to decompose the causes of the chains' front-store revenue collapse. The results attribute the collapse to the changed shopping environment: rising distance costs shifted demand away from the dedicated trips on which the format depends, and this channel alone accounts for more than half of the decline. The evolving configuration of competing stores actually worked in the chains' favor, and their relative price increases, evaluated against remaining demand, were roughly revenue-neutral.

On the supply side, a static pricing game recovers marginal costs and markups by inverting Bertrand–Nash first-order conditions under the national, firm-level pricing that the major retailers practice. This converts the demand estimates into store-level variable profits in dollars. A dynamic game then models entry and exit at the level of individual retail sites. Because flow profits enter the model in dollars from the demand system, the dollar scale of the unobserved payoff shocks is an estimable parameter rather than a normalization. A finite-dependence argument allows estimation without solving the game (Arcidiacono and Miller, 2011). Assuming that exit is terminal and a vacated site returns to the entrant pool, every value beyond the period differences out of the stay/exit comparison, leaving estimating equations that are linear in the cost parameters and require no value function, equilibrium fixed point, or stance on firms' beliefs about the future reimbursement path.

The dynamic cost parameters face the classic endogeneity of comparing profits with exit. Under the free-entry assumption, profitable markets have high rents, so regressing exit on variable profit would suffer from correlation between fixed costs and variable profit. In addition to the cross-section understating how profit moves exit, the aggregate time series confounds the drug margin decline with everything else that happened to the industry over the years. Entry costs are therefore identified from a within-county-year contrast between entrants and incumbents that are otherwise under similar market conditions. The dollar scale is then identified from a shift-share design: state-year changes in Medicaid fee schedules and Medicare Part D remuneration shift the dispensing margin differentially across counties and, within a county, differentially across stores in proportion to their fill volumes—variation in profit that originates from payers rather than with the store's own market conditions.

I then check cost structure against the industry reports. The implied fixed operating cost of a chain store, \$3.8 million per store-year, is bracketed by the per-store operating costs reported in Walgreens' financial statements: \$3.5 million in fiscal 2023 and \$4.1 million in fiscal 2024. The independent pharmacy estimate of \$1.25 million lies within the range observed in proprietary multi-state independent-pharmacy income statements and in the NCPA Digest, though toward its upper end. Because a reimbursement counterfactual shifts variable profit only, its computation is invariant to the cost levels and requires only the dollar scale, the discount factor, and the estimated choice probabilities (Aguirregabiria and Suzuki, 2014).

I use my estimated model to run counterfactuals highlighting how restoring the dispensing margin to its real-terms peak—a \$4.17 increase per fill, the strongest version of any reform on the table—would reshape the pharmacy network. The no-policy baseline continues to decline by a striking 26,000 stores, while the long-run national pharmacy count rises by only 9.3 percent relative to it. The induced stores locate overwhelmingly where fills are plentiful, which is where pharmacies already operate: 98.1 percent of the induced store-mass accrues to counties with three or more pharmacies, while counties with at most one pharmacy gain only 16 stores nationwide, as they were already operating in a thin market and the profit gain is not sufficient for a new entry. Because the higher margin is paid on every fill, the policy costs \$30.1 billion per year, of which only 0.8 percent is paid at stores whose presence it causes. The availability gain is worth about one cent to consumers per public dollar spent.

The findings run counter to the alleged pharmacy desert claims and policy proposals. The reimbursement lever is real but economically insignificant because the subsidy pays every fill everywhere. At the same time, the closure problem is concentrated in a format whose shortfall lies in retail revenue that reimbursement does not touch. In fact, even the long-run projected closures under the current market conditions would not create access problems. Exits concentrate in dense markets, and drug-focused formats backfill where access binds. If anything, the relevant risk is the viability of the traditional chain business model, not proximity for drug-fillers.

Related Literature

This paper contributes to four strands of literature. A large body of public health research has coined the term "pharmacy deserts" and examined their consequences for adherence and health (Qato et al., 2014, 2017; Amstislavski et al., 2012; Qato et al., 2019; Guadamuz et al., 2021a,b; Wittenauer et al., 2022, 2024; Ying et al., 2022). Recently, quasi-experimental estimates using econometric methods study what closures do to patients (Caputi, 2026; Battles, 2026) and review the pharmacy channel's role in healthcare delivery (Mulligan, 2023). These works measure where pharmacies are and what happens when they disappear, but take the store configuration itself as given. They also provide primarily reduced-form evidence and only speculation on the causes. I model the configuration as the equilibrium outcome of entry, exit, and pricing decisions. This allows me to ask whether closures harm patients and whether counterfactual reimbursement policies would change where pharmacies operate and, if so, how much that would benefit patients.

Second, the paper connects to work on the changing retail landscape such as the diffusion of big-box retail and its effects on incumbents (Jia, 2008; Holmes, 2011; Arcidiacono et al., 2020), dollar-store and e-commerce expansion (Leung and Li, 2022; Hortacsu and Syverson, 2015), and the consequences of retail access for consumption, as in the food-desert debate (Atkin et al., 2018; Allcott et al., 2019). The pharmacy market offers a distinctive case study of the changing retail landscape, in which the store format bundles an essential good whose price is neither controlled by the retailer nor the consumer. The setting allows me to study how general-merchandise retailing spills over into healthcare access through the store's other categories.

Third, the supply side builds on the literature on spatial competition and dynamic entry and exit (Seim, 2006; Houde, 2012; Sweeting, 2013; Aguirregabiria and Suzuki, 2014; Bodéré, 2023; Almagro and Domínguez-Iino, 2025; Schneier, 2025). Embedding dollar-denominated flow profits from an estimated demand system into the dynamic problem is standard in this literature. The distinguishing feature of the paper is how the scale of payoffs is identified: I exploit an observed profit shifter that is excluded from the demand side. In particular, statutory changes in dispensing fees move per-fill margins. At the same time, the fill volume is inelastic, and consumer prices are unchanged, allowing the shifter to trace the mapping from dollars to entry and exit behavior without passing through the demand system. The shift-share structure, where state-level fee changes are interacted with predetermined payer mix, delivers site-level variation. The estimator exploits the finite-dependence structure of the stopping problem (Hotz and Miller, 1993; Arcidiacono and Miller, 2011), with counterfactuals resting on the identification results of Aguirregabiria and Suzuki (2014) and

Kalouptsidi et al. (2021).

Fourth, the demand model contributes to the literature on multi-category demand, one-stop shopping, and supermarket pricing (Hess and Gerstner, 1987; Bell and Lattin, 1998; Smith, 2004; Song and Chintagunta, 2007; Gentzkow, 2007; Rhodes, 2015; Thomassen et al., 2017; Dubé et al., 2025; Cao et al., 2024). The literature established that cross-category effects discipline multi-product retail pricing, with the loss leader as the leading example. The pharmacy setting differs from the classic loss-leader good as a third party prices drugs. Therefore, the retailer cannot choose the depth of its loss leader, and the complementarity that sustains the format can erode for reasons outside the retailer's control. My paper separates the causes of this erosion and measures how that pricing friction affects the market structure.

The rest of the paper is organized as follows. Section 2 describes the institutional setting and the data. Section 3 documents the stylized facts on closures, margins, and the drug–CPG linkage. Section 4 develops the multi-category demand model, Section 5 its identification and estimation, and Section 6 the demand estimates, model fit, and the decomposition of the front-store decline. Section 7 embeds the demand system in the supply side, where a static pricing game recovers marginal costs and variable profits, and a dynamic game over retail sites models entry and exit. Section 8 reports the counterfactual policy experiment, and Section 9 concludes.

2 Institutional Setting and Data

In this section, I define the part of the pharmacy industry under study and describe the data that allow me to measure demand, margins, market structure, and adherence outcomes. I focus on retail, outpatient-facing pharmacies that are community locations that dispense prescription and over-the-counter medications to the public. The focus of the paper, therefore, excludes inpatient hospitals, closed-door long-term care facilities, and specialty or central-fill pharmacies that are not open to the public. Big-box and regional grocery stores that operate in-store pharmacy counters fall within the scope because they are also community-facing outlets.

Within retail outpatient pharmacies, I distinguish two business models that differ in product variety and typical scale: large, multi-product chain-store pharmacies ("chains") that sell prescription drugs (Rx) alongside a broad assortment of consumer packaged goods (CPG),[1] and smaller, independent pure-play pharmacies ("independents") that focus on dispensing and carry limited general merchandise. The supply model refines this split into five permanent operator types—chain pharmacies, independent pharmacies, big-box or grocery stores with pharmacy counters, grocers without pharmacies, and smaller CPG-only formats such as

[1] Throughout the paper, I refer to *CPG* for the full non-prescription retail assortment, which encompasses fresh groceries, packaged groceries, and other household essentials all together. When the non-grocery component alone is meant, I call it *other household essentials*; it is one of the three retail categories of the demand model (Section 4).

dollar and convenience stores (Section 7.2). Stores are classified by national chain names, complemented by the NAICS code classification otherwise, with each company assigned its modal type to eliminate spurious type churn. Appendix A.3 provides more detail about the classification.

I work with several complementary levels of aggregation, each serving the analysis's objective. I estimate demand at the consumer–shopping-occasion level on a national random sample of households and shopping days (Section 5). The dynamic model of entry and exit tracks individual retail sites observed annually from 2013–2024, with the reimbursement margin measured at the county-year level (Section 7). The population-level distributions of distance and driving time to the nearest pharmacy summarize access (Section 3). I estimate drug adherence responses to closures on event-time panels centered on the month of a local exit. Entry and exit of stores are measured from the establishment panel described below, and I reconcile name and ownership changes so that rebrandings and ownership transfers are not counted as closures.

The empirical analysis exploits multiple data sources, each serving a distinct role: consumer demand (Numerator); retailer locations and pharmacy availability (Infogroup (Data Axle), complemented by NCPDP); reimbursement rates and prescription-drug utilization (Merative MarketScan, NCPA); payer-side inputs and policies to the reimbursement (NADAC, CMS); and external validation of the supply model's dollar-denominated cost estimates (a proprietary multi-state independent pharmacy operator). The primary estimation dataset for demand is the **Numerator household panel**, a consumer-level panel of shopping trips that records retailer and store identifiers, dates, item descriptions, quantities, and prices, along with rich household demographics.[2] Most importantly, Numerator covers both prescription fills and CPG purchases within the same trip basket. It therefore supports trip-level complementarity measures across product categories, as well as store-level demand, category price indices (Appendix A.2), and foot-traffic measures that are used in demand estimation. I estimate the demand model of Section 4 on this panel, and I describe the estimation sample in Section 5 and the category definitions in Appendix A.1. One limitation of the Numerator data is that it records only whether a prescription was purchased, not the drug-level details that characterize prescription composition or adherence — such as the molecule, the National Drug Code (NDC), or days supplied.[3]

Retailer locations, entry, and exit come from the **Infogroup Historical Business Registry**,[4] an annual longitudinal map of U.S. retail establishments with addresses, geocodes, industry codes, and store attributes including store size. I construct the store-year establishment panel for both descriptive facts and the supply model, and I use it to link store attributes to Numerator store lists. The activity spells are built to measure

[2] Numerator provides exact store coordinates along with the company names, which makes the dataset more suitable for this paper than other syndicated consumer panel datasets such as NielsenIQ Marketing Data.

[3] However, this is still better than other syndicated consumer panel datasets; NielsenIQ Marketing Data, for instance, censors any prescription data purchases.

[4] Now called Data Axle.

economic entry and exit rather than data coverage, and ownership transfers are bridged so that they are not counted as closures. Appendix A.3 details the construction.

Reimbursement rates and prescription drug fills come from two complementary sources. **Merative MarketScan** commercial claims records include allowed total amounts paid to the pharmacy, patient cost sharing, and NDC-level detail, including days supplied. I use it to measure plan-side payments and out-of-pocket trends, build standard adherence measures, and estimate closure-event studies, as described in Section 3. The **National Community Pharmacists Association (NCPA) Digest** provides an annual national per-prescription margin series for independent pharmacies, which serves as the benchmark level of dispensing profitability throughout the paper. It serves as the anchor for the county-year margin series described below and the basis for the real-terms peak that defines the headline counterfactual (Section 8.2).

Two additional sources complement drug costs and pharmacy location counts. The **National Average Drug Acquisition Cost (NADAC)** series reports a monthly, NDC-level proxy for pharmacy purchase prices. The spread between MarketScan allowed amounts and NADAC yields the pre-fill gross margin for prescription. The **NCPDP dataQ** pharmacy master file lists licensed dispensing pharmacies with addresses, store attributes including pharmacy types, and closure dates, and is used to validate the registry-based pharmacy universe and the timing of closures.

Public data from the **Centers for Medicare & Medicaid Services (CMS)** provide payer-side inputs and policies governing reimbursement rates for the supply model. Combining CMS enrollment and fee-schedule data with the NCPA benchmark yields the county-year dispensing margin. This serves as a state variable and as the basis for identifying variation in the dynamic cost parameters. Section 7 and Appendix A.4 detail the construction.

Finally, the estimated fixed costs for independent pharmacies are validated through a data agreement with a **multi-state independent pharmacy operator**, whose identity is confidential. I observe transaction-level dispensing records from more than eighty pharmacies across eight Midwestern states from mid-2022 onward, including their drug acquisition costs, payer reimbursements, quantities, cash prices (if purchased with cash), and patient characteristics, as well as store-level income statements for 2023 and 2024. The income statements directly report observed per-store operating costs including pharmacist salaries and rents, against which I check the dollar-denominated fixed costs recovered by the supply model (Section 7.7).

I apply the following conventions throughout the paper. I adjust all dollar-denominated magnitudes to 2024 dollars using a CPI-U index, the appropriate common unit for comparing profits and policy costs across years (Section 7.3). I also introduce any derived measures where they are used: the complementarity and access measures in Section 3, the estimation sample and instruments in Section 5, the variable-profit construction in Section 7.3, and the county-level drug margin variation in Section 7.

3 Stylized Facts on Retail Pharmacy Closures

This section demonstrates that the retail pharmacy industry has been contracting and that closures disrupt the patients they displace. I then argue that the primary driver of the decline is the erosion of CPG traffic and co-purchasing in locations with abundant retail competition, rather than declining drug margins. Two features of the disruption from closures matter for the subsequent sections: it is *temporary*, because patients relocate their fills rather than abandon them, and it is *spatially concentrated* in places dense with surviving pharmacies, so that distance and driving time to the nearest pharmacy are essentially unaffected.

3.1 Industry Is Declining, but Closures Have Little Lasting Effect on Patients

Fact 1: Net closures are concentrated among chains. The pharmacy industry expanded until 2021 and has contracted since (Figure 1). The net closures of major pharmacy chains, CVS, Walgreens, and Rite Aid, rather than of independents, primarily drove this reduction. The major-chain count falls steadily while independent and retailer (big-box and grocery) pharmacy counts rise over the overall contraction period (Figure 2).

Fact 2: Closures temporarily reduce local drug adherence. When a pharmacy closes, the patients it displaces experience a decline in adherence. Still, the decline is temporary, and the recovery, as well as pre-period behavior, is as informative as the dip itself. I show this with two complementary exercises: a *patient-level* event study that follows displaced patients from pharmacies that closed, and a *market-level* panel regression that asks whether markets experiencing a net loss of pharmacies fill fewer prescriptions. Both point to the same conclusion: patients continue to fill their medications after closure.

Patient-level evidence: the closure event study. I follow patients of closing pharmacies and ask whether their own prescription supply falls, replicating the closure-event studies of Caputi (2026) and Battles (2026) but using the MarketScan commercial claims data. Both of those studies work with Medicare claims, so their evidence pertains to an older population, arguably the one more harmed by closures and most relevant for policy, as they tend to be high fillers. The replication therefore offers a check on their findings on a different demographic group. The unit of treatment is a pharmacy identifier that no longer appears. In MarketScan, the identifier distinguishes pharmacy–insurance, so a "closure" is the disappearance of a pharmacy–insurance cell rather than necessarily the closure of a physical pharmacy.[5] I designate a pharmacy

[5]The MarketScan pharmacy identifier is defined by a data contributor and is frequently missing, so a vanishing cell can reflect a genuine closure, a network or PBM change that drops the pharmacy, or the contributor's exit from the panel. The continuous-enrollment restriction below removes the last case. The first two remain, making the treatment a noisier access shock than a physical closure would be. Hence, a small intensive-margin effect is a conservative finding rather than an artifact of clean treatment.

as closing in the month of its last observed claim, provided it serves at least thirty distinct patients over its life. It is followed by at least twelve months of zero claims inside the sample, which distinguishes closure from right-censoring. I assign a patient to a closing pharmacy if more than half of her fills in the three months before closure occurred at that pharmacy (Caputi, 2026).

I estimate a stacked Poisson event study (Cengiz et al., 2019). For each closure cohort, defined as a calendar month of closure, the treated group is the assigned patients and the control group is patients whose own pharmacy closes at least twelve months later, so that they are not yet treated during the event period (Deshpande and Li, 2019). Stacking the cohorts and indexing patient–month observations by event time k relative to closure,

$$\mathbb{E}\big[y_{ict} \mid \text{treated}_i, k_{ct}\big] = \exp\Big(\sum_{k \neq -1} \beta_k \, \mathbf{1}\{\text{treated}_i\} \cdot \mathbf{1}\{k_{ct} = k\} + \alpha_{ic} + \lambda_{ct}\Big), \tag{1}$$

where y_{ict} is patient i's days of prescription supply in month t within cohort c. The patient-by-cohort effects α_{ic} absorb patient composition, and the cohort-by-calendar-month effects λ_{ct} absorb the common event-time path, so the β_k are identified from the treated–control contrast with the month before closure ($k = -1$) omitted. The Poisson form retains patient-months with zero fills, and the β_k are proportional effects on days supplied. The sample is restricted to patients continuously enrolled with drug coverage by an insurance over the entire window. This is especially important for analyzing commercial claims, as it purges the mechanical decline that would otherwise arise when a data contributor exits the panel, leaving genuine within-enrollment access shocks. Standard errors are clustered on the pharmacy identifier, and the outcome counts *all* of a patient's fills-at any pharmacy, including fills with a missing pharmacy identifier, and through mail order. Doing so ensures that relocating to another pharmacy or switching to mail delivery registers as a continued supply rather than a loss.

Figure 3 reports the estimates. The coefficients are flat from twelve to four months before closure, supporting the parallel-trends assumption; the final three pre-closure months show a fill surge of roughly five percent, consistent with patients (and the closing pharmacy) pulling fills forward ahead of the closure. The post-closure effect is sharp but short-lived. Days supplied fall by roughly ten percent relative to the month before closure (about five percent relative to the earlier pre-closure path), with the trough in the first three months and a return to the pre-closure path within roughly six months. These magnitudes are in line with what Caputi (2026) and Battles (2026) find in the Medicare population, where the response can be further decomposed (fewer retail trips, offset by 90-day and mail-order fills). In both Medicare patients and commercial patients, the estimated effects are not economically significant. The disruption is a dip of a few percentage points in days supplied that closes within months, in both the population where fill burdens are heaviest and the one studied here. For this paper, the central finding is that patients continue to obtain their medications by relocating fills, which is why the adherence effect is transitory rather than permanent.

I exploit the observed temporary nature of the disruption as an input to the demand model. Section 4.4 imposes a deterministic-fill regime, where a consumer with a drug need fills with certainty, and so a closure changes *where* the fill occurs, never *whether* it occurs. That regime would be untenable if closures caused permanent losses in prescription consumption, since a persistent post-closure decline in adherence would be evidence that fills are destroyed rather than relocated. The event-study recovery and pre-period pulling show that fills are relocated at the patient level, and the market-level evidence below shows the same at the aggregate level.

Market-level evidence: no decline in prescription consumption. Whether the transitory disruption aggregates into a market-level decline in prescription consumption is a related but separate question, perhaps with more policy relevance. A reimbursement subsidy is justified on access grounds only if the contraction in pharmacy availability reduces the number of prescription days patients fill. I therefore ask whether patients of markets that lose pharmacies *on net* consume fewer prescription days.

I merge the Infogroup pharmacy panel to MarketScan claims at the metropolitan-area (MSA) level, the geographical aggregation level of MarketScan, and estimate a within-patient fill response,

$$\text{DaysSupply}_{it} = \beta \, \Delta \log \text{Stock}_{m(i),t} + \eta_i + \tau_t + \varepsilon_{it},$$

where $\Delta \log \text{Stock}_{m,t}$ is the *net* annual change in the count of active pharmacies in patient i's MSA, and η_i, τ_t are patient and year fixed effects, with standard errors clustered at the MSA level. Identification comes from within-patient, over-time variation in own-market availability, so β can be interpreted as the elasticity of prescription consumption with respect to local pharmacy availability. I report the Poisson estimate, which delivers β directly as an elasticity and retains enrolled patients with zero fills, keeping the extensive margin in the outcome.

The elasticity of local pharmacy availability as reported in Table 1 is a precise zero. The point estimate implies that a ten percent net contraction in the local pharmacy stock moves prescription days by only 0.03 percent of the mean, and the confidence interval rules out losses larger than 0.4 percent of days from a contraction of that size.

Fact 3: Exits concentrate where pharmacies are dense, leaving distance and driving time to the nearest pharmacy unchanged. Closures cluster in places already served by many pharmacies. Between 2021 and 2024, counties that had a single pharmacy in 2021 *gained* stores on net, and net losses rise steeply with the 2021 pharmacy count: counties with more than fifty pharmacies account for 58 percent of all net losses (Figures 4 and 5). This pattern is inconsistent with a uniform drug-revenue shock, which would have hit all markets similarly. In fact. because exits concentrate where surviving pharmacies are close by, the population-level distributions of distance and driving time to the nearest pharmacy are essentially unchanged between 2021 and 2024 (Figures 6 and 7), and the prevalence of pharmacy deserts has not risen. The industry has lost redundancy in dense markets, not coverage at the margin where coverage matters for access.

3.2 The Cause: Declining Drug-CPG Complementarity, Not Drug Margins

Drug margins have fallen, and consumers are not paying more out of pocket. Prescription drug margins have declined over the sample period (Figures 8 and 9), which helps explain why policymakers have cited drug profitability as the cause for closures. At the same time, patient out-of-pocket payments, including copays, coinsurances, and deductibles, have remained roughly stable (Figure 10).

Fact 4: CPG sales explain exits more than falling drug margins do. The CPG expenditure share at chain pharmacies has fallen (Figure 12), the share of drug trips that include a CPG purchase has fallen (Figure 13), and total CPG spending at chain pharmacies has declined (Figure 11) even as chains raised CPG prices.

I conduct a reduced-form exercise that tests the two candidate drivers to ask whether the closures track the front-store decline or the margin decline. From the establishment panel, I measure each county's net loss of pharmacies between adjacent years and relate it to one-year-lagged county-level measures of the two channels. For the front store, I use the chain pharmacies' share of the county's CPG expenditure and the share of chain pharmacy trips on which a prescription and CPG are purchased together (both from the household panel, in percentage points). For dispensing, I use the county-year margin of Section 2, expressed as a percentage of the average prescription charge. The estimating equation is a fixed-effects Poisson for the net-loss *rate*—net losses relative to the county's incumbent stock—over 2019–2024:

$$\mathbb{E}\big[\text{NetLoss}_{mt} \mid x_{m,t-1}\big] \;=\; n_{m,t-1}\, \exp\Big(\beta_1\, \text{CPGShare}_{m,t-1}+\beta_2\, \text{DualShare}_{m,t-1}+\beta_3\, \text{DrugMargin}_{m,t-1}+\eta_m+\tau_t\Big), \tag{2}$$

where n_{mt} is the county's active pharmacy count and the lagged count $n_{m,t-1}$ enters as an offset. The county fixed effects η_m absorb permanent market characteristics, the year effects τ_t absorb the national trend, and the standard errors are clustered at the state level. Identification is within-county, where each coefficient compares the same county's loss rate in years when the lagged regressor was above versus below its county mean, so permanent differences across counties, such as size, demographics, and retail structure, do not contribute.

Table 2 reports the estimates. The front-store channel predicts closures: a one-percentage-point decline in the chain pharmacies' share of a county's CPG spending raises the county's net-loss rate by roughly 14 percent in the following year. The dual-purchase share adds essentially nothing beyond it. The dispensing margin has no predictive power as the point estimate is small relative to its standard error and, if anything, wrong-signed. Even though the dispensing margin is the lever that every reform proposal focuses on, pharmacies disappear where and when the front store erodes, not where and when dispensing pays less.[6]

This regression is a natural first pass at the question, and its limitations motivate the model that follows. First, the regressors are equilibrium outcomes, not primitives. The chain pharmacy CPG share moves when consumers reallocate their shopping, but also mechanically when a pharmacy closes, which would be the reverse of the causality the regression is after. And the margin coefficient is the least reliable, but that might be because reimbursement policy responds to distress, with higher-fee policies implemented precisely because pharmacies close, which would make the coefficient more likely to find that higher margins accompany more closures.

The deeper limitation is that the reduced-form evidence could not answer the question the paper poses: *why* pharmacies are closing and *whether* and *how* the proposed policies would salvage consumer welfare. Suppose I directly obtain each pharmacy's profit, decomposed by category, in dollars. That accounting would show that both the front store and drug profits eroded, and by how much, but not why. Is it because consumers' preferences shifted, because the set of competing stores changed, or because of the prices the chains themselves set? This distinction affects the welfare question. If consumers stopped valuing what the chain front store sells, the closures would be the market's efficient contraction, and consumers would lose little from exits they have already left. In this case, a reimbursement subsidy would only pump up capacity that shoppers no longer value. If, instead, the chains' own prices pushed trips away, the same closures would carry a different welfare weight and warrant different policy remedies. A county-level association cannot separate these explanations, and no reduced-form coefficient can say how many stores a restored margin would retain, where, or what they would be worth to the consumers. Those questions require primitives on preferences and costs that stay fixed while prices, choice sets, and policies change one at a time, and recovering them is the task of the model discussed in the next sections.

[6] In fact, the correlation between the market size and the pharmacy type composition is worth a note. Figure 14 shows that small markets with fewer pharmacies tend to comprise independent pharmacies. How this correlation contributes to the closure pattern is later discussed in Section 7.

4 A Model of Multi-Category Retail Demand

The stylized facts attribute retail pharmacy exits to the erosion of CPG demand in pharmacies, driven by declines in convenience foot traffic and by prescription fillers deciding not to purchase CPGs. Evaluating this explanation and the reimbursement policies motivated by the competing drug-margin narrative requires a demand system in which consumers jointly choose where to shop among stores that offer different subsets of categories, how much to buy in each category, and whether to fill a prescription. The demand model aims to deliver store-level demand by category, the cross-category response of CPG purchases to a prescription fill, and cross-store substitution when a store closes or changes its prices.

The model rests on quadratic utility over category quantities—the benchmark specification for demand systems with binding non-negativity constraints since Wales and Woodland (1983)—and is embedded in the multi-category, multi-store shopping framework that Thomassen et al. (2017) developed to study supermarket pricing. Consumers facing shopping costs visit one or two stores per occasion and choose continuous quantities in several product categories. The quadratic utility makes conditional demands linear at interior solutions and makes corner solutions well-defined, and its off-diagonal curvature measures cross-category complementarity in preferences. In fact, these curvature terms serve as an object distinct from the correlation in category tastes that panel data can separately identify (Gentzkow, 2007). Shopping costs prompt one-stop shopping, so demand in one category responds to the prices of other categories sold at the same store. Once internalized by a multi-category retailer, these cross-category effects discipline its prices, and the same internalization logic is the economic engine of the chain pharmacy business model: prescriptions generate trips, and trips are monetized in high-margin retail categories.

I adapt Thomassen et al. (2017)'s framework to fit my setting. I model three continuous categories chosen to track the format distinctions at the center of the closure story—fresh groceries (f), packaged groceries (p), and other household essentials (x)—and let category availability differ across formats. Major chain pharmacies carry no fresh groceries, whereas grocery stores and big-box retailers carry all three. Splitting groceries into fresh and packaged components matters because the retailers whose expansion most plausibly eroded pharmacy CPG traffic (big-box stores and local grocery stores with pharmacy counters) differ from pharmacies precisely in their fresh assortments.

I add prescription drugs as a fourth, *binary* category ($y \in \{0, 1\}$). Payers, rather than stores, set drug prices, so there is no usable cross-sectional variation. The drug therefore enters demand through store-level availability and through its complementarity with the continuous categories, not through a price term. I treat filling as the consequence of a latent health need rather than a discretionary purchase. Conditional on a need, the consumer fills with certainty, and the model determines *where* the fill occurs and how it shifts the rest of the basket (Section 4.4). I also introduce pharmacy-type intercepts in drug utility that absorb payer steering of fill location across pharmacies. Then, fill-conditional purchase *quantities* for CPGs rather than where consumers fill identify the complementarity parameters.

4.1 Setup

A shopping occasion is a consumer-day (i, t). On each occasion, consumer i faces a location-specific set of stores j, each belonging to a format (banner type) $b(j)$—chain pharmacy, independent pharmacy, big-box, grocery, dollar store, etc. There are four categories of consumption:

1. Fresh groceries (*Fresh*), a continuous quantity $q_{ic}^{f} \in \mathbb{R}_{+}$;
2. Packaged groceries (*Packaged*), a continuous quantity $q_{ic}^{p} \in \mathbb{R}_{+}$;
3. Other household essentials (*Other*), a continuous quantity $q_{ic}^{x} \in \mathbb{R}_{+}$;
4. Prescription drugs (*Drug*), a binary fill decision $y_{ic} \in \{0, 1\}$.

Store j posts prices p_j^k for the continuous categories $k \in \{x, f, p\}$ it carries. These prices are *endogenous* choices of the retailer.[7] Payers set the prescription price, which is common to all pharmacies facing a given consumer, and displays no usable cross-sectional variation. Hence, the drug carries no price term in demand.

Availability differs across stores in two ways, both of which are artifacts of the data rather than parameters. First, a category can be absent from a store's assortment: chain pharmacies (CVS, Walgreens, Rite Aid) carry no fresh groceries, so the (j, f) price is missing, and fresh groceries cannot be purchased there. Second, a prescription can be filled only if the consumer has a drug need ($R_i = 1$) and a chosen store operates a pharmacy ($S_c = 1$), imposing the feasibility constraint

$$y_{ic} \leq R_i S_c. \tag{3}$$

Consumers may spread a shopping trip across multiple stores. A *basket* c is a set of one or two stores, and $\mathcal{C}_{it}$ denotes the full feasible set of one- and two-store baskets in i's choice set. Two features of the data make this structure essentially without loss: consumers visit at most two stores on roughly 92 percent of occasions, and they buy each category at a single store within the basket on roughly 96 percent of occasions (Table 5). The remaining corner cases in which a consumer visits a store but buys zero units of category k are generated by the model itself via the non-negativity constraints below. The consumer buys each continuous category at the member store offering the highest net-of-price value for that category. The within-basket

[7] Empirically, p_j^k is a banner-by-week hedonic price index estimated from the household panel's item-level transactions. Appendix A.2 details the construction, the treatment of missing store-category cells, and the associated price instrument.

store assignment is therefore mechanical given prices and tastes, and the category is forced to zero if no store in the basket carries k. A basket remains in the choice set even when it cannot supply every category. For instance, a consumer could buy household essentials, packaged groceries, and fill a prescription at a chain pharmacy. The consumer buys no fresh groceries, but, in fact, no fresh groceries are available at chain pharmacies, and the availability structure exists to rationalize such trips.

4.2 Preferences

Conditional on the store choice c and the category-store assignment, let

$$A_{ic}^{k} = \max_{j \in c\,:\,k \text{ available}} \left(\mu_{itjk} - \alpha_i \, p_j^k\right), \qquad k \in \{x, f, p\},$$

denote the net-of-price value of category k, where μ_{itjk} is i's taste for category k at store j (specified in Section 4.6) and α_i is the marginal utility of income. Writing $q_{ic} = (q_{ic}^x, q_{ic}^f, q_{ic}^p)'$, the variable utility from the basket is

$$V_i(c, q_{ic}, y_{ic}) = \sum_{k \in \{x,f,p\}} A_{ic}^k \, q_{ic}^k \;-\; \tfrac{1}{2}\, q_{ic}' \, \Lambda \, q_{ic} \;-\; \left(\lambda_{xy} q_{ic}^x + \lambda_{fy} q_{ic}^f + \lambda_{py} q_{ic}^p\right) y_{ic} \;+\; \kappa_c \, y_{ic}, \tag{4}$$

subject to $q_{ic} \geq 0$, the availability restrictions, and (3), where

$$\Lambda = \begin{pmatrix} \lambda_{xx} & \lambda_{xf} & \lambda_{xp} \\ \lambda_{xf} & \lambda_{ff} & \lambda_{fp} \\ \lambda_{xp} & \lambda_{fp} & \lambda_{pp} \end{pmatrix} \succ 0. \tag{5}$$

Utility is quasilinear in an outside good. The specification is not an ad hoc functional form: (4) is the reduced form of a fully general store-and-quantity choice problem under four restrictions—the two-store, one-store-per-category shopping structure read off the data, additively separable shopping costs, quasilinearity in the outside good, and weak separability with homothetic within-category subutilities. Appendix B.1 discusses the restrictions and what recommends the modeling choice of quadratic utility over familiar alternatives.

Λ is a positive-definite matrix common to all consumers: the diagonal terms $\lambda_{kk} > 0$ govern own-category satiation, and the off-diagonal terms $\lambda_{xf}, \lambda_{xp}, \lambda_{fp}$ capture intrinsic cross-category effects among the continuous categories (negative values denote complements). Because drug is binary, $y_{ic}^2 = y_{ic}$, so an own-drug quadratic term would be linear in y_{ic} and indistinguishable from the linear drug value. The drug therefore carries no own-curvature parameter and interacts with the quadratic system only through the cross terms $\lambda_{xy}, \lambda_{fy}, \lambda_{py}$. When these are negative, filling a prescription *raises* the marginal utility of the continuous categories, which would imply the drug–retail complementarity.

$$\kappa_c = \max_{j \in c\,:\,S_j=1} \mu_{j,y}$$

is the drug value of the basket's best pharmacy. A needer who chooses basket c fills at the member pharmacy with the highest drug appeal $\mu_{j,y}$, and that value is credited to the basket.

4.3 Conditional Demands and Corner Solutions

Fix a basket c and a feasible fill decision y. The quantity problem is a concave quadratic program, and its first-order necessary conditions[8] are the Kuhn–Tucker system

$$\Lambda\, q^*_{ic} \;\geq\; a_{ic}(y), \qquad q^*_{ic} \;\geq\; 0, \qquad q^{*\prime}_{ic}\big(\Lambda\, q^*_{ic} - a_{ic}(y)\big) = 0, \tag{6}$$

where $a_{ic}(y) = \big(A^x_{ic} - \lambda_{xy} y,\ A^f_{ic} - \lambda_{fy} y,\ A^p_{ic} - \lambda_{py} y\big)'$ collects the fill-shifted net values and unavailable categories are fixed at zero. (6) holds with equality for the categories purchased in positive quantity, and conditional demands are *linear* in the net values:

$$q^{k*}_{ic} = \frac{1}{\lambda_{kk}}\Big(A^k_{ic} - \lambda_{ky}\, y - \sum_{k' \neq k} \lambda_{kk'}\, q^{k'*}_{ic}\Big) \qquad \text{for each purchased category } k, \tag{7}$$

while for a category at a corner, the marginal value of the first unit is non-positive net of the interaction terms. The diagonal λ_{kk} scales the response of category k's quantity to its own price, and the continuous off-diagonals $\lambda_{kk'}$ transmit purchases of k' into the marginal value of k, and a fill shifts every purchased category's demand through λ_{ky}.[9]

Because the objective is strictly concave, the optimum is unique and can be computed exactly by enumerating active sets. Note that corner solutions are model predictions rather than nuisances. They rationalize the zero-purchase outcomes that are pervasive in the data, including trips to stores where fresh groceries are unavailable by construction.

[8] The positive definiteness of Λ makes it sufficient as well

[9] The complementarity parameters also have a quantity interpretation. In the one-continuous-category special case where we have a single composite good CPG other than drugs, the (interior) change in CPG demand induced by filling a prescription at the same store is

$$\Delta q^{x*} \equiv q^{x*}_{ic}(y=1) - q^{x*}_{ic}(y=0) = -\frac{\lambda_{xy}}{\lambda_{xx}}, \tag{8}$$

so a negative λ_{xy} implies that consumers buy more CPGs on fill trips, holding the store fixed. In the full model, the same object is the fill-induced shift of the vector $a_{ic}(y)$ through the corner-aware demand map.

The indirect (deterministic) value of the basket c is

$$w_i(c) = \max_{y\,:\,y \le R_i S_c} \Big\{ a_{ic}(y)' q^*_{ic} - \tfrac{1}{2}\, q^{*\prime}_{ic} \Lambda\, q^*_{ic} + \kappa_c\, y \Big\}, \tag{9}$$

evaluated at the corner-aware quantities. Therefore, predicted quantities, purchase incidence, and the discrete store choice are mutually consistent.

Store-level demand. The firms price at the national level,[10] so the relevant demand object aggregates over baskets. Demand for category k at store j is

$$Q_{jk} = \sum_{c \,\ni\, j} \mathbf{1}\{c \text{ chosen}\}\, \mathbf{1}\{j = d_k(c)\}\; q^{k*}_{ic}, \tag{10}$$

where $d_k(c)$ denotes the within-basket category assignment. A change in store j's price of category k moves (10) on three margins: an *intensive* margin, the change in q^{k*} holding the basket and assignment fixed, governed by (7); a *reassignment* margin, the change in which member store supplies k within a two-store basket; and an *extensive* margin, the change in the basket choice itself. Cross-category effects at the store level arise on the extensive margin even under separable preferences, where a consumer who abandons store j over its essentials prices takes the prescription fill along—and on the intensive margin through the interaction terms. The demand derivatives entering the supply-side pricing conditions of Section 7 aggregate all three margins, which is why the recovered markups internalize the complementarity.[11]

4.4 Prescription Drugs: Latent Need and Deterministic Fill

Prescription filling differs economically from the continuous categories. For most patients on most occasions, whether to fill is dictated by a health need rather than by comparison shopping, whereas *where* to fill and what else to buy on the trip are genuine choices. I encode this asymmetry by taking the limit in which the net value of filling grows large ($\kappa_r \to \infty$): a consumer with a drug need fills with certainty, and a consumer without a need never fills.

Under this assumption, the observed trip-level fill indicator directly reveals the need type R_i. The limit is well-behaved since it operates on the choice set rather than on the estimating equations. A needer collects the fill value in every admissible basket—pharmacy-containing baskets, and only those—so κ_r enters each admissible alternative as a common constant and cancels from the basket-choice probabilities (11) at any

[10] For instance, CVS sets a single price for a given week for packaged groceries

[11] The assignment $d_k(c)$ depends on prices only through the linear index $\mu_{itjk} - \alpha_i p^k_j$, not on the chosen quantities, because Λ carries no store-specific elements. This is what makes the assignment mechanical given the basket.

value. The limit removes pharmacy-free baskets from the needer's choice set and fixes $y = 1$.[12]

The regime is grounded in the reduced-form evidence rather than assumed for convenience. Fact 2 of Section 3 shows that patients displaced by a closure suffer only a *temporary* dip in days supplied. Fills are relocated, not destroyed, and market-level prescription consumption does not respond to net pharmacy availability (Table 1). A permanent post-closure drop in filling would be direct evidence against the regime, but the data instead show exactly the pattern the regime implies—a change in where fills occur with no change in whether they occur.

- **Non-needers** ($y_{it} = 0$): No fill anywhere. Continuous demands solve (7) at $y = 0$.
- **Needers** ($y_{it} = 1$): Only pharmacy-containing store combinations remain in the store choice set. The fill occurs at the basket's best pharmacy. Continuous demands solve (7) at the complementarity-shifted vector $a_{ic}(1)$, and the basket value adds κ_c.

This deterministic-fill regime determines what the drug side of the model can and cannot identify. Need prevalence is pinned to the observed fill rate rather than estimated. I cannot identify the *level* of drug utility because the value of the fill itself is present in every feasible store choice of a needer. However, I can still identify cross-category complementarities $\lambda_{xy}, \lambda_{fy}, \lambda_{py}$, identified from the gap in continuous purchases between fill and non-fill occasions. I can also identify the set of *differences* in pharmacy drug intercepts $\xi_{b(j),y}$, which govern where the fill lands.

Because the drug carries no taste shocks or size gradient, the drug appeal of store j is $\mu_{j,y} = \xi_{b(j),y}$ exactly, and $\kappa_c = \max_{j \in c: S_j = 1} \xi_{b(j),y}$. Only differences across firms matter, as a common shift cancels in the needer's choice probabilities, so I normalize the reference format, independent pharmacies, to zero.[13] Because the out-of-pocket prescription price is set by plan design rather than by the pharmacy, the drug enters without a price term, and these intercepts absorb the *net-of-copay* appeal of filling at each format relative to an independent pharmacy.[14] The intercept terms' role in the identification strategy is to separate fill *location* from fill-conditional *quantities*. Without them, payer steering of fills toward particular formats would have to be rationalized by the complementarity and shopping-cost parameters, contaminating the parameters of interest (Section 5.7).

[12] Therefore, no unbounded quantity enters any moment, and what remains are finite utility differences across baskets, formats, and fill states.

[13] Firms that never operate pharmacies yield structural zeros.

[14] This may include preferred-network copay differentials, 90-day-at-retail programs, and pharmacist relationships. The absorption is exact for any out-of-pocket differential that is constant at the format level, and such a differential is observationally equivalent to a drug-format taste, so the constant-price treatment does not require that out-of-pocket prices be literally identical across pharmacies. Appendix B.3 states the weaker condition it requires and reports the supporting claims evidence.

4.5 Store Choice

Each occasion involves nested choices:

$$\underbrace{\text{basket } c \in \mathcal{C}_{it} \ (\leq 2 \text{ stores})}_{\text{discrete}} \rightarrow \underbrace{\text{store-per-category assignment}}_{\text{mechanical given } c} \rightarrow \underbrace{\text{quantities } q_{ic}, \text{ fill } y_{ic}}_{\text{continuous + incidence}}.$$

Adding a basket-level taste shock ε_{ic}, drawn i.i.d. Type-I extreme value with unit scale, and integrating it out, the basket choice is a logit *conditional on the taste draws* ν (Section 4.6):

$$P_i(c \mid \nu) = \frac{\exp\big[w_i(c) - \Gamma_i(c)\big]}{\sum_{c' \in \mathcal{C}_{it}} \exp\big[w_i(c') - \Gamma_i(c')\big]}, \tag{11}$$

where for needers, $\mathcal{C}_{it}$ contains only pharmacy-including baskets. Unlike a textbook logit, the "mean utility" $w_i(c)$ has no closed form: it is the value of a corner-aware quadratic program, recomputed for every draw of ν. The shopping cost $\Gamma_i(c)$ that includes distance enters only at this discrete layer and is excluded from variable utility. This exclusion restriction separates shopping-cost parameters from variable-utility parameters.

4.6 Consumer Heterogeneity

Tastes. The taste of consumer i for category k at store j on occasion t combines a format-by-category fixed effect with a scaled common index and four Gaussian taste shocks:

$$\mu_{itjk} = \xi_{b(j),k} + \beta_{0k}\Big(\underbrace{\beta_h \,\mathrm{hh}_i + \beta_s \ln(\mathrm{size}_j) + \psi_t}_{\text{common index}} + \underbrace{\sigma_1\nu_i + \sigma_2\nu_{it} + \sigma_3\nu_{ik} + \sigma_4\nu_{ijk}}_{\text{taste shocks}}\Big), \tag{12}$$

where hh_i is household size, size_j is store floor space (in square feet), ψ_t is a time (year-quarter) effect, and all $\nu \sim N(0,1)$. The fixed effects $\xi_{b(j),k}$ carry the levels. The category scales $\beta_{0k} > 0$ let a common index and common shocks load differently on each category, with the normalization $\beta_{0x} \equiv 1$ anchoring units to the essentials category.

The four shocks vary across levels, and their persistence is the identifying structure that the panel exploits. ν_i is a permanent consumer-level shopping-intensity shock common to all categories. ν_{it} is the only *time-varying*, occasion-level shock common across categories within a trip. ν_{ik} is a permanent consumer-category taste. ν_{ijk} is a permanent consumer-store-category taste capturing horizontal store differentiation and loyalty. The common shocks $\sigma_1\nu_i + \sigma_2\nu_{it}$ generate *correlation* in category tastes, which must be separated from genuine complementarity in Λ. For the drug, $\beta_{0,y} \equiv 0$, so drug appeal reduces to the format intercepts of Section 4.4.[15]

Price sensitivity. Price sensitivity includes an income-dependent mean and a multiplicative Rayleigh shock:[16]

$$\alpha_i = \left(\alpha_1 + \frac{\alpha_2}{\text{income}_i/\text{hh}_i}\right)\nu_i^\alpha, \qquad \nu_i^\alpha \sim \text{Rayleigh}(1), \tag{13}$$

so that price sensitivity falls with income per household member and $\alpha_i > 0$ for all i whenever $\alpha_1, \alpha_2 > 0$.

Shopping cost. The shopping cost includes a random fixed cost of visiting a second store with a random per-distance cost:

$$\Gamma_i(c) = (\gamma_{11} + \gamma_{12}\nu_{i1}^\Gamma)\,\mathbf{1}[n(c) = 2] + (\gamma_{21} + \gamma_{22}\nu_{i2}^\Gamma)\,\text{dist}_{ic}, \qquad \nu_{i1}^\Gamma, \nu_{i2}^\Gamma \sim N(0,1), \tag{14}$$

with $n(c)$ the number of stores in the basket and $\text{dist}_{ic} = 2\sum_{j \in c} \text{dist}_{ij}$ the round-trip distance. The spreads γ_{12}, γ_{22} are permanent consumer-level heterogeneity in two-stop and travel costs (car access, time constraints, rurality).

The model contains two conceptually distinct sources of cross-category linkage—structural complementarity through the off-diagonal λ's, and the shopping-cost channel, by which costly trips concentrate purchases at the store that serves as the anchor. The model identifies them separately: the λ's from fill-conditional quantity gaps within stores and the shopping-cost parameters from discrete basket choice. Appendix B.2 further discusses the two complementarity concepts and the detailed crosswalk to Thomassen et al. (2017).

5 Identification and Estimation of the Demand Model

The full parameter vector is

$$\Theta = \underbrace{\{\beta_{0f}, \beta_{0p},\ \xi_{bk},\ \beta_h, \beta_s,\ \psi,\ \alpha_1, \alpha_2,\ \Lambda,\ \lambda_{xy}, \lambda_{fy}, \lambda_{py}\}}_{\text{means and curvatures}} \cup \underbrace{\{\sigma_1, \ldots, \sigma_4,\ \gamma_{12}, \gamma_{22}\}}_{\text{spreads}} \cup \{\gamma_{11}, \gamma_{21}\} \cup \{\xi_{b,y}\}.$$

Under the deterministic-fill regime, the need prevalence is pinned to the data, the drug-value level and scale are held, and the drug-side unknowns reduce to the complementarities $\lambda_{xy}, \lambda_{fy}, \lambda_{py}$ and the format-intercept differences $\xi_{b,y}$. I estimate Θ by the simulated method of moments.

[15] The normalization of the basket shock completes the model. Rescaling every utility parameter, including α_i, Λ, and the shopping costs, together with the shock scale, leaves all choice probabilities and all quantities unchanged. But then only ratios to the shock scale are identified. The anchor $\beta_{0x} \equiv 1$ pins *relative* category loadings, not the overall scale. Dollar-denominated objects are invariant to the convention because they always divide by α_i.

[16] This shock keeps the price sensitivity term positive.

5.1 Sample and Descriptive Statistics

The consumer panel has too many shopping trips, making demand estimation on this sample computationally expensive. I therefore estimate it on a random sample of the national panel: 3,000 households, observed on a random sample of days drawn from each quarter of 2019–2024. I deliberately sample days at random rather than using each household's complete daily history. Shopping trips are serially dependent, in which a trip yesterday replenishes inventories and affects today's shopping decision. The model treats occasions as independent draws conditional on the permanent taste heterogeneity, and randomly spaced days make the occasion-level shock plausibly i.i.d. across a household's sampled occasions. This lets the cross-period moments read persistence off permanent heterogeneity rather than off short-run trip chaining. The screening steps that impose the basket structure of Section 4 retain over 90 percent of expenditure (Table 4). I discuss the remaining sample construction choices in Appendix B.4.

Table 3 reports the sampled households' demographics, the choice environment they face—the availability of, distance to, and size and price level of each retailer within thirty miles—and the dispersion of the category price indices on which the price parameters are identified. Table 5 reports the shopping outcomes that motivate the model's structure. Consumers visit a single store on 68 percent of occasions and at most two stores on 92 percent of occasions. Conditional on visiting two stores, only 12 percent of a category's spending occurs at the second store. In 4

5.2 Selection, Corners, and the Role of the First-Order Conditions

I now discuss the challenges an estimator for this class of models faces, and how their resolution shapes the empirical strategy (Thomassen et al., 2017, Section III). First, consumers choose which stores to visit, and the same parameters and unobserved tastes drive both store choices and quantities purchased at those stores. An analysis using only chosen stores alone would therefore confound preferences with selection into stores for the quantity choices. I address this by estimating all parameters simultaneously from the store choice indicator and the quantities conditional on it. Doing so forces the selection rule to be part of the estimator rather than a nuisance, and the exclusion restriction introduced in Section 4 disciplines this joint estimation. Also, while distance strongly shifts basket choice, I exclude it from the variable utility, giving the discrete layer its own identifying variation, in the manner of exclusion restrictions in selection models.

I also find prevalent zero purchases. However, a visit to a store with no fresh-grocery purchases could be an equilibrium outcome of the model rather than a missing observation. While the common approach is to append a censoring device to an interior demand system, I work directly with the first-order conditions (6). For any candidate parameter vector Θ, taste draw ν, basket, and fill state, the Kuhn–Tucker system is solved exactly by enumerating the candidate active sets[17]. The restricted linear system (7) delivers closed-form quantities, and the utility-maximizing non-negative candidate is the unique global optimum. This computation produces the predicted quantities q^*, the purchase indicators $\mathbf{1}\{q^{k*} > 0\}$, and the indirect value $w_i(c)$ that feeds the basket-choice probabilities (11). The first-order conditions therefore drive the estimation, and predicted outcomes are by-products of optimality conditions evaluated at candidate parameters. This way, the estimator requires no share inversion, no fixed point, and no numerical search within the inner problem.[18]

The structural assumptions that impose the model's outcomes at the true parameters are exactly the assumptions a maximum-likelihood estimator would require with taste draws independent of observables. Indeed, the model's likelihood exists (its supermarket analog is derived in Appendix G of Thomassen et al., 2017). But evaluating that likelihood requires computing the probability of the observed basket-and-quantity outcome, and the corner structure makes it computationally prohibitive. Instead, moment conditions built on prediction errors require only conditional means and remain consistent for a fixed number of draws (McFadden, 1989; Pakes and Pollard, 1989). The approach thereby joins the discrete–continuous demand tradition that models store and quantity choices jointly (Smith, 2004; Song and Chintagunta, 2007; Dubé et al., 2025) with the older literature on demand systems under binding non-negativity constraints (Wales and Woodland, 1983).

5.3 Estimation Principle: Simulated Prediction Errors

For any observed outcome Y^* generated by the model, $Y^* = Y(\Theta_0, x, \nu, \varepsilon)$ with $(\nu, \varepsilon) \perp x$,

$$\mathbb{E}\big[\, Y^* - Y(\Theta_0, x) \,\big|\, x \big] = 0, \qquad Y(\Theta, x) = \int Y(\Theta, x, \nu, \varepsilon)\, f(\nu, \varepsilon)\, d(\nu, \varepsilon). \tag{15}$$

The extreme-value shocks ε are integrated analytically (the logit P_i), and the taste draws ν are integrated by simulation over $R = 25$ draws per consumer, held fixed throughout the optimization. Because the fill itself is deterministic given the need, the fill-incidence moment is identically zero for any parameter value and is therefore pruned. Fill-conditional quantities discipline the drug side and fill location instead. The within-period predictions are sample means over the simulation draws; writing $P_c^{(r)}$, $q^{k*(r)}$, and $d_k^{(r)}(c)$ for the basket probability, the conditional quantities, and the within-basket store assignment for category k

[17] Restricting store choice to 2 stores and working with 3 continuous categories yields at most 2^3.

[18] The cost, relative to a textbook discrete-choice model, is that $w_i(c)$ has no closed form and must be recomputed, as the value of a small quadratic program, for every basket and every draw. The same machinery that is evaluated once at the estimated parameters generates the predictions used for model fit and the decomposition exercises of Section 6, as well as the demand derivatives consumed by the supply side.

evaluated at draw $\nu^{(r)}$,

$$\widehat{I}_c = \frac{1}{R}\sum_{r=1}^{R} P_c^{(r)}, \quad \widehat{Q}_{cjk} = \frac{1}{R}\sum_{r=1}^{R} q^{k*(r)}\,\mathbf{1}\{j = d_k^{(r)}(c)\}\,P_c^{(r)}, \quad \widehat{D}_{cjk} = \frac{1}{R}\sum_{r=1}^{R} \mathbf{1}\{q^{k*(r)} > 0\}\,\mathbf{1}\{j = d_k^{(r)}(c)\}\,P_c^{(r)},$$

for each continuous category $k \in \{x, f, p\}$.

5.4 Within-Period Moments

For each consumer-occasion, I match observed to predicted quantities, purchase incidence, and the basket choice, interacting the prediction errors with instruments Z and stacking the three blocks, summed over occasions t, baskets $c \in \mathcal{C}_{it}$, and member stores $j \in c$:

$$g_i^{(1)}(\Theta) = \begin{bmatrix} \sum_{t,c,j} Z_{cjk}^{Q}\big(Q_{cjk}^{*} - \widehat{Q}_{cjk}\big) \\ \sum_{t,c,j} Z_{cjk}^{D}\big(D_{cjk}^{*} - \widehat{D}_{cjk}\big) \\ \sum_{t,c} Z_c^{I}\big(I_c^{*} - \widehat{I}_c\big) \end{bmatrix}, \qquad k \in \{x, f, p\}, \tag{16}$$

where starred objects are their data analogs and the Q and D rows stack over the three continuous categories. The instruments are category-separated. Z^Q contains a constant, demographics (income, household size, age), log store size, the own price, the two cross prices of the other continuous categories, a two-store dummy, the *fill indicator* $\mathbf{1}\{y_{it} = 1\}$, format fixed effects, and time dummies; Z^D is Z^Q without the time dummies; and Z^I contains distance, the two-store dummy, distance squared, their interaction, the mean price of the basket, the mean price divided by income, and *fill-by-pharmacy-format* columns that compare the observed fill location with the model's. The last columns constitute the fill-location block: among needers, the share of fills landing at each pharmacy format, matched to the choice-probability-weighted format composition of the model's predicted fill stores.

The demographics, store size, store format, and time interactions in Q and D pin β_{0k}, ξ_{bk}, β_h, β_s, and ψ_t; own- and cross-price interactions pin Λ; price levels in Q, D and the price terms in Z^I pin α_1, α_2; the fill-indicator columns in Q, D pin the complementarities $\lambda_{xy}, \lambda_{fy}, \lambda_{py}$ through the fill-conditional shift of continuous demand; the fill-location columns in Z^I pin the format-intercept differences $\xi_{b,y}$; and the geometry columns in Z^I pin the shopping-cost means γ_{11}, γ_{21}.

Price variation identifies α and the diagonal curvature on *different margins*, exactly as the first-order conditions suggest. λ_{kk} is pinned by how much less of k consumers buy when its price rises, the intensive margin operating in the Q and D moments through (7). α_1, α_2 are pinned by how much less likely consumers are to include the store in the basket at all—the extensive margin operating in the I moment.

5.5 Cross-Period Moments

The cross-sectional moments cannot distinguish the taste-shock spreads from one another because the four shocks are observationally similar within a single occasion. The panel separates them through persistence. Define the model objects per occasion and per draw (all weighted by the basket choice probabilities P_c)

$$r = \sum_{c,j,k} p_j^k q_{cj}^k P_c, \quad q_k = \sum_{c,j} q_{cj}^k P_c, \quad d_{jk} = \sum_{c \ni j} \mathbf{1}\{q_{cj}^k > 0\} P_c,$$

$$os = \sum_c \mathbf{1}[n(c) = 1] P_c, \quad dist = \sum_c \text{dist}_{ic} P_c, \quad v_j = \sum_{c \ni j} P_c,$$

that is, total spending, category quantities, store-category purchase incidence, one-stop frequency, expected distance, and store-visit probability. The cross-period moments match *products across adjacent occasions*, for example $R(\Theta) = \int r(\nu_t)\, r(\nu_{t-1})\, f\, d\nu$ against $r_t^* r_{t-1}^*$:

$$g_i^{(2)} = \sum_{t \geq 2} \Big[R^* - R,\ \textstyle\sum_k (Q_k^* - Q_k),\ \sum_{j,k} (D_{jk}^* - D_{jk}),\ OS^* - OS,\ DIST^* - DIST,\ VP^* - VP \Big]. \quad (17)$$

Adjacent occasions share the permanent draws $(\nu_i, \nu_{ik}, \nu_{ijk}, \nu^\alpha, \nu^\Gamma)$ but draw the occasion shock ν_{it} independently, so cross-period covariance flows *only* through the time-invariant heterogeneity. Cross-period covariance in total spending pins σ_1; in category quantities, σ_3; in store-category incidence, σ_4; and in one-stop frequency and distance, the shopping-cost spreads γ_{12}, γ_{22}.

5.6 Cross-Category Moments

The remaining confound is the pair (σ_1, σ_2). Both shocks are common across categories and generate cross-category co-movement that could masquerade as complementarity when it is actually due to correlated tastes. Following the logic of Gentzkow (2007), they are separated by comparing *within-period* and *cross-period* cross-category spending covariances. With category spending $r_k(\nu) = \sum_{c,j} p_j^k q_{cj}^k P_c$, define for each pair $k < k'$

$$R_{kk'}^{\text{in}} = \int \tfrac{1}{2} \big[r_k(\nu_{t-1}) r_{k'}(\nu_{t-1}) + r_k(\nu_t) r_{k'}(\nu_t) \big]\, f\, d\nu \qquad \text{(within period)},$$

$$R_{kk'}^{\text{cr}} = \int \tfrac{1}{2} \big[r_k(\nu_{t-1}) r_{k'}(\nu_t) + r_k(\nu_t) r_{k'}(\nu_{t-1}) \big]\, f\, d\nu \qquad \text{(across periods)},$$

and match both to their data analogs:

$$g_i^{(3)} = \sum_{t\geq 2} \left[\, R_{kk'}^{\text{in}*} - R_{kk'}^{\text{in}}, \quad R_{kk'}^{\text{cr}*} - R_{kk'}^{\text{cr}} \,\right] \qquad \text{for all pairs } (x,f),\ (x,p),\ (f,p). \tag{18}$$

The occasion-level shock $\sigma_2 \nu_{it}$ correlates two categories' spending within an occasion but not across occasions, whereas the permanent shock $\sigma_1 \nu_i$ does both; the within-versus-cross gap identifies σ_2 separately from σ_1. With three continuous categories, each cross-curvature $\lambda_{xf}, \lambda_{xp}, \lambda_{fp}$ requires its own pair to be disentangled from taste correlation, which is why all three pairs enter.

5.7 Identifying the Drug Interactions

Drug has no price variation, so $\lambda_{xy}, \lambda_{fy}, \lambda_{py}$ cannot use the cross-price experiment that identifies the continuous cross-effects. What that experiment requires, however, is not a price as such but an exogenous shifter of one category whose effect on the others can be traced, and the drug has one: the health need itself. Conditional on the taste draws, a need switches the fill from zero to one and shifts the retail first-order conditions (6) by exactly the interaction vector, so λ_{ky} is read off how the same consumer's retail quantities and purchase incidence differ between fill and non-fill occasions at the same store—the fill-indicator columns of the Q and D moments.

The identifying condition is an exclusion restriction. The fill indicator shifts retail demand only through the drug need R_i and pharmacy availability S_c, neither of which enters retail demand directly. As in Gentzkow (2007), the panel allows me to separate this shift from taste correlation, since there is no price experiment. A consumer's non-fill occasions measure her permanent shopping intensity, category tastes, and store loyalties $(\sigma_1, \sigma_3, \sigma_4)$, the within- versus cross-period covariances separate the occasion-level common shock (σ_2), and the interactions are what remain of the fill-conditional gap once those channels are absorbed.

The threat to this restriction is that CPG-carrying stores bundle a pharmacy with other retail selection, so pharmacy availability correlates with unobserved store quality. Four features of the design address the concern. First, format-by-category fixed effects mean the fill-conditional comparisons are made *within* store formats, differencing out the average quality gap between pharmacy-carrying and other stores. Second, the panel compares the same consumer's retail basket on fill and non-fill occasions at the same store, so time-invariant store quality cannot account for the gap. Third, the object of interest is the *decline* in the complementarity over time, and a stable availability-quality confound cannot produce it. Fourth, the pharmacy-format drug intercepts absorb fill-specific store appeal that need *not* be stable over time. With sorting absorbed there, the λ's are pinned by fill-conditional quantities rather than by where fills locate.

5.8 Price Endogeneity

Endogeneity concerns apply only to the retail prices p^x, p^f, p^p, since the drug price is payer-set.[19] CPG prices are measured as firm-level hedonic indices constructed from item-level transactions (Appendix A.2), consistent with the national pricing policies of the major retailer chains during the sample period. The baseline absorbs market-level unobserved quality with format-by-category fixed effects ξ_{bk} as well as time dummies, so that $(\nu, \varepsilon) \perp \text{price} \mid x$. Therefore, the baseline rests on the conditional-independence assumption rather than on instruments.[20]

5.9 Estimator and Inference

Stacking the blocks $g_i = (g_i^{(1)\prime}, g_i^{(2)\prime}, g_i^{(3)\prime})'$ and averaging across consumers, the estimator is

$$\widehat{\Theta} = \arg\min_{\Theta} \ \bar{g}(\Theta)' W^{-1} \bar{g}(\Theta), \qquad \bar{g} = \tfrac{1}{N} \textstyle\sum_i g_i, \tag{19}$$

with moments clustered by consumer. The weighting matrix here is the inverse of the covariance matrix $W = N^{-1} \bar{g}(\tilde{\Theta})' \bar{g}(\tilde{\Theta})$ and $\tilde{\Theta}$ are first-stage estimates.

6 Demand Estimates and Model Fit

6.1 Interpretation

Table 6 reports the demand estimates. I estimate every parameter of the full random-coefficients model separately for each year's occasions, with only the winsorization caps and centering constants frozen at their pooled values to ensure comparability across columns. Nothing in the construction ties one column to the next, so if the model is recovering stable preference primitives rather than fitting year-specific noise, the columns should agree.[21] Moreover, if the collapse of the front store reflected a drift in the primitives themselves, it should appear here as trending coefficients.

[19] The drug-price treatment has exogeneity and retailer-invariance assumptions. *Exogeneity*: out-of-pocket price is determined by plan design rather than chosen by the pharmacy in response to demand conditions. This is why I impose that the drug carries no price moment and no analog of the instruments below. *Retailer-invariance*: a given consumer faces the same out-of-pocket price wherever the fill lands. It further justifies omitting the drug price from utility altogether.

[20] As a safeguard, I construct Hausman-style instruments for each firm-category-week, where I take the leave-one-out average of the other banners' indices (Appendix A.2) (Hausman et al., 1994). Instrument validity requires that cross-banner price co-movement within a week reflect common costs rather than common demand shocks, with the aggregate component of the latter absorbed by the time dummies. Re-estimating with these instruments in the price moments is a planned robustness exercise.

[21] A more formal way to tackle this is to allow for time-varying coefficients and estimate the specification on the full-year sample. This would allow for direct tests on the time component of the coefficients and could work with a fixed set of firm-category fixed effects, rather than testing their invariance as well. However, full-sample demand estimation is computationally prohibitive, and the year-specific demand estimates do not exhibit meaningful anomalies that warrant the fully specified version.

Tastes are stable over time, but price sensitivity rose in real terms. The taste block, the utility curvature, and the heterogeneity spread change little over six years, including pre- and post-pandemic years. The category loadings β_{0f} and β_{0p} oscillate without trend (1.45–1.54 and 0.46–0.50), the own-curvatures $\lambda_{xx}, \lambda_{ff}, \lambda_{pp}$ stay within narrow bands, and the taste spreads σ_1–σ_4 are nearly constant. The largest movement anywhere in the taste and price blocks is the price coefficient α_1, which rises from 0.32 to 0.37 in 2021 and reverts by 2024—and because the price indices are nominal, this movement reads differently in real terms. A flat nominal α_1 across a period of 23 percent cumulative inflation is a *rising* marginal utility of the constant dollar. Converted to 2024 dollars, price sensitivity jumps by roughly a fifth with the 2021 inflation surge and stays there, so the apparent reversion is a nominal artifact.[22] In contrast, the dollar itself bought steadily less, and the marginal utility of *income* rose. This suggests that as inflation compressed real budgets, an additional constant dollar came to matter more, and consumers emerged from the high-inflation era more price-conscious than when they entered it.

Year-to-year differences elsewhere are small relative to their standard errors and show no monotone drift. Six independently estimated columns that agree when there is no economic reason for them to differ indicate that the specification is recovering primitives rather than fitting year-specific noise. Moreover, the decline in front-store activity is not a broad shift in what consumers want from stores. In fact, the one primitive in the taste and price blocks that moved is the value of a dollar, which is not a taste for any store or category but a sharpening of the price margin on which every store competes; together with the costlier distance documented below, it is the shift in the shopping environment that the decomposition of Section 6.3 finds responsible for the front-store decline.

The drug–retail interactions run against taste-based complementarity, and the purchase linkage lives in shopping costs. The drug interactions enter as $-\lambda_{ky} q^k y$, so complementarity in preferences requires $\lambda_{ky} < 0$, and a positive estimate means that a fill *lowers* same-trip demand for category k, holding the store and tastes fixed (equation (8)). The estimates are positive: the essentials–drug term λ_{xy} sits at roughly 0.5 in every year and is statistically distinguishable from zero throughout; the packaged-grocery–drug term λ_{py} is of comparable magnitude (0.41–0.45), and the fresh-grocery–drug term is indistinguishable from zero. Conditional on the store visited and the taste draws, a fill occasion therefore carries a somewhat *smaller* storable-goods basket. Fill trips resemble errands, not occasions on which the marginal utility of the front store is elevated. The sign is stable, and it sharpens the interpretation of the drug–retail complementarity that sustains the chain business model. The complementarity is a property of *store-level demand*, generated by the shopping-cost channel, not a taste primitive that ties the prescription to the pharmacy's CPGs.

[22] The conversion affects only this parameter because quantities are deflated expenditures; all other estimates are already in real quantity, distance, or utility units.

The reading has a precedent in Thomassen et al. (2017), which finds that estimated utility interactions across supermarket categories are likewise small and, where signed, are substitutes. They also agree that shopping costs generate strong one-stop shopping. Preferences are therefore not a driver for co-purchasing. Instead, the basket is captured by the trip, which is why supermarkets, big-box stores, and online channels can contest it. There is also no trend: λ_{xy} moves from 0.531 in 2019 to 0.498 in 2024, a difference well inside a single standard error. Hence, a drift in the preference interaction is not what changed over the period, despite the observed decline in the drug-CPG co-purchasing behavior.

Distance has become costlier, but a second stop has not. The clearest time variation in the table is in the shopping costs, and it is asymmetric. The per-distance cost γ_{21} rises from 0.136 in 2019 to 0.170 in 2024, a 25 percent increase and about two and a half standard errors of the difference. At the same time, the fixed cost of visiting a second store, γ_{11}, is flat to slightly declining (4.63 to 4.32), and the spread γ_{22} is unchanged. Consumers have not become more reluctant to coordinate multi-stop trips; rather, the cost of covering distance has risen. Converted at α, the increase carries through to nominal dollars per mile, but in constant 2024 dollars, the per-mile cost is roughly flat. Both distance and money became more consequential relative to the occasion shocks. This asymmetry implies where the lost front-store spending went. A rising distance cost penalizes every basket in proportion to the travel it requires. It penalizes dedicated trips the most, while leaving the cost of adding a stop or a category at a store already being visited unchanged. It therefore favors channels that require no travel at all, consistent with the observed decline of CPG spending coinciding with the increased spending at online retailers.

Bottom line. The model admits the preference interaction λ_{xy} channels and the shopping-cost channel through which a prescription fill can generate retail demand. The yearly demand estimates show that the preference interaction is essentially flat over 2019–2024, implying that a weakening drug–CPG linkage in preferences does not explain the front-store decline. What changed is the shopping environment—the real value of money and the utility cost of distance—and the decomposition below finds that this shift alone more than accounts for the front-store revenue decline, with the evolution of the choice set partially offsetting it and the chains' relative price increases roughly revenue-neutral against the demand that remained.

6.2 Model Fit

I assess the model fit based on its ability to reproduce the shopping patterns that matter for the supply side, both in the estimation sample and out-of-sample on data not used in estimation. Table 7 reports, year by year, the correlation between predicted and observed store-category demands, for both quantities and shopper counts (Panel A), together with the mean absolute error in firm shares of category demand (Panel B1) and in the one-stop share of category demand (Panel B2). Table 8 disaggregates to the firm level, comparing predicted and observed revenue and shopper shares across the major reta. Table 9 does the same for the drug side, comparing the predicted and observed composition of fill locations.

The dimensions the supply side uses are the dimensions the model fits best: the correlations between predicted and observed store-category demands exceed 0.97 in every year, out of sample as well as in, and firm shares of category demand are matched to within about a percentage point on average—Walmart's roughly 20 percent revenue share, the pharmacy banners' single-digit shares, and the long tail of regional formats are all reproduced closely (Table 8). There is essentially no deterioration from the in-sample to the out-of-sample columns, so the fit does not reflect overfitting of the estimation sample. While the model reproduces the broad ranking of fill locations, it underpredicts the share of fills for chain pharmacies (roughly 20 percent versus 29 percent for CVS). It overpredicts fills at big-box and food-and-drug formats. This is the margin the pharmacy-format drug intercepts, and the fill-location moments are built to absorb it, so the remaining gap indicates that format-level intercepts do not capture banner-specific steering, such as preferred networks concentrated within particular chains. Throughout, the out-of-sample columns evaluate the model on a holdout of 3,000 additional households drawn from the panel and processed identically to the estimation sample.

6.3 Decomposing the Decline in the Chain Front-Store Business

The demand estimates locate the change of the period not in tastes for drug–CPG linkage but in the shopping environment. What, then, drives the collapse of the chain front-store business? The estimated model can answer this question because it prices any combination of primitives, observed or counterfactual. Therefore, I decompose chain pharmacies' real front-store revenue into spending across retail categories, excluding prescription fills and the dispensing margin, which falls on the supply side of Section 7. Between the model's 2021 and 2024 baselines, this revenue falls by 26.6 percent in 2024 dollars, and the exercise replicates that decline one channel at a time.

I start with the 2021 market, when the pharmacy count was the highest. I begin with the 2021 choice sets, 2021 prices, and the 2021 demand parameters, and move to the 2024 market in sequence: (i) *preferences*, replacing the 2021 demand parameters with their 2024 values, which asks whether behavior changed because what consumers want changed; (ii) *availability*, replacing the 2021 configuration of stores available to each household with the 2024 configuration; and (iii) *prices*, replacing 2021 retail prices with 2024 prices. The demand solves underlying the decomposition are inflation-free.[23] Figure 15 reports the resulting path, and

[23] In mixed-year cells, the price coefficient is rescaled by the CPI ratio, so the price channel measures *relative*-price movements

the three steps sum exactly to the total decline.

Replacing the 2021 demand parameters with their 2024 values in the otherwise unchanged 2021 market removes roughly 1.5 times the total decline. Its content is the shift in the shopping environment documented in Section 6.1. In particular, the rising cost of distance penalizes the trips that chain pharmacies depend on. It favors channels that require no travel, rather than a decline in the taste for what pharmacies sell or a drift in the drug-interaction primitives. The jump in real price sensitivity is already reflected in the 2021 baseline, so it enters the decomposition not through the parameter changes but as the elasticity used to evaluate the other channels. The availability runs in the opposite direction: replacing the 2021 store configuration with the 2024 configuration *raises* chain front-store revenue, recovering about 44 percent of the decline. The net evolution of the choice set over this period—rival closures alongside the chains' own consolidation—left the spatial distribution of stores *more* favorable to the chains, not less. Lastly, the final step to 2024 relative prices is essentially revenue-neutral. The decomposition, therefore, reads the collapse of the front store as a shift in demand away from the format. The changed shopping environment more than accounts for the decline on its own, while the evolution of the store network partially offset it.

7 Supply: Pricing, Costs, and the Dynamics of Entry and Exit

I now embed the estimated demand system into the dynamic oligopoly model of pharmacy entry and exit. First, I recover marginal costs and markups for the retail categories assuming a static pricing game. Here, I assume that firms set prices for the categories they control (fresh groceries, packaged groceries, and other household essentials). Payers set prescription drug prices, failing to satisfy any first-order pricing condition for the retailers, and I use the county-year dispensing margin as the reimbursement lever. Second, I use the recovered margins and the model-implied quantities to derive store-level variable profits that encompass both profit channels. Third, I model a dynamic game over retail *sites* that incumbents can occupy, potential entrants can claim, and exiting stores can vacate. The dollar scale of the unobserved payoff shocks is an estimable parameter rather than a normalization with flow profits entering from the demand system. A finite-dependence argument allows estimation without fully solving the game.

7.1 Static Pricing and Marginal-Cost Recovery

Retail prices are set at the level of the retail firm, consistent with the national pricing policies documented in DellaVigna and Gentzkow (2019). Firm F chooses the prices of the continuous categories it carries, A_F,

rather than the mechanical passage of inflation, and the real-terms evolution of price sensitivity travels with the demand parameters in channel (i).

to maximize variable profit,

$$\Pi_F = \sum_{\ell \in A_F} \left(p_\ell^F - mc_\ell^F\right) Q_\ell^F(\Theta) \; + \; m^r \, Y_F(\Theta),$$

where Q_ℓ^F is the demand for firm F's category-ℓ offering implied by the estimated model, m^r is the per-script drug margin, and Y_F is the firm's prescription volume, with the payer-set drug price taken as given. The first-order condition for each $\ell \in A_F$ is

$$Q_\ell^F + \sum_{\ell' \in A_F} (p_{\ell'}^F - mc_{\ell'}^F) \frac{\partial Q_{\ell'}^F}{\partial p_\ell^F} + m^r \frac{\partial Y_F}{\partial p_\ell^F} = 0.$$

Stacking across $\ell \in A_F$ and inverting yields the vector of firm–category marginal costs,

$$mc^F = p^F + \Delta_F^{-1}\Big(Q^F + m^r \frac{\partial Y_F}{\partial p^F}\Big), \qquad [\Delta_F]_{\ell\ell'} = \frac{\partial Q_{\ell'}^F}{\partial p_\ell^F}, \tag{20}$$

where $p^F = (p_\ell^F)_{\ell \in A_F}$, $mc^F = (mc_\ell^F)_{\ell \in A_F}$, and Δ_F collects the within-firm demand derivatives. All demand objects and derivatives come directly from the estimated demand model—the derivatives by finite differences of the simulated demand system—so no additional instruments are needed.

The final term in (20) is the prescription-internalization correction. Because every needer fills every occasion when there is a need, aggregate prescription volume is fixed and $\sum_F \partial Y_F/\partial p_\ell = 0$. That is, raising a retail price does not eliminate fills; it diverts them to rival chains. A chain with a large fill volume and a strong trip complementarity therefore prices its retail categories below the single-product benchmark to protect prescription traffic. In contrast, independents, which sell little retail, carry corrections near zero.

I rely on parametric bootstrap over the estimated distribution of the demand parameter to quantify the sampling uncertainty in the recovered costs. Figures 16 and 17 report the recovered marginal costs and the implied profit margins as distributions across store–product categories, comparing 2019 with 2024, with pointwise 90 percent bootstrap bands. The cost distribution shifts to the right in every category over the period, while the margin distribution compresses toward zero. Chain pharmacies shift from the middle of the 2019 margin distribution to the bottom by 2024, while the online format sustains the highest margins throughout.

7.2 The Dynamic Environment: Types, Sites, and Timing

Store types. Markets are counties, indexed by m, observed annually over $t = 2013, \cdots, 2024$ in the establishment panel. Each retail establishment j carries a permanent type $\tau \in \mathcal{T} = \{C, I, B, G, P\}$: chain pharmacies ($C$), independent pharmacies (I), big-box stores with pharmacies (B), grocers without pharmacies (G), and CPG-only formats such as dollar and convenience stores (P). I assume that there is no format switching: an incumbent of type τ chooses only between continuing as τ and exiting, consistent with industry practice over the sample period.[24] I formalize the categorization process in Section 7.2.

Entry and exit are strategic for the pharmacy types, $\mathcal{T}^{\text{en}} = \{C, I\}$. The remaining types are non-strategic in two distinct cases. The dollar-and-convenience type (P) shares the small-box site universe with the pharmacy types and transitions through estimated choice probabilities of the same form as theirs—a type-specific exit probability and a branch of the site-level entry model—but these are reduced-form descriptions of its arrival and survival, not structural choices. Therefore, no entry cost is attached, and the counterfactuals hold the P probabilities at their baseline values.

I do not explicitly model the evolution of the large-box types (B, G) and treat them as exogenous state variables. The literature documents that their location decisions are national portfolio choices made as part of broader format strategies (Holmes, 2011) and, by that, would be invariant to the local pharmacy configuration and to dispensing margins.[25] That is, they enter the model only as covariates that shift pharmacy profits and choice probabilities, measured locally via distance to the nearest big-box store and counts within distance bands around a site, rather than as county totals. I therefore freeze the big-box locations at the last observed configuration for counterfactual simulations (Appendix C.5).

Site as the unit of analysis. The unit of analysis for the strategic types is the physical *site*—an individual retail location—rather than a market-level count of stores.[26] I make this deliberate model choice as counterfactuals become direct. A counterfactual state *is* a configuration of pinned locations, so the demand model of Section 4 runs without any auxiliary mapping from counts back to store locations.

The site universe. For county m, the site set $\mathcal{S}_m$ collects every location ever occupied during 2013–2024 by a small-box format—chain pharmacy (C), independent pharmacy (I), or CPG-only store (P)—in the establishment panel of Appendix A.3. I define a site as a spatial cell of roughly 0.1 miles around the establishment coordinates, motivated by the observation that same-format store co-location does not occur within this range, even in the most dense retail areas such as Manhattan. Each site s carries an integer capacity q_s, the maximum number of establishments observed operating concurrently at the cell (measured

[24] While there have been recent announcements of major chain pharmacies opening small-format drug-only stores, their entry is after the sample period; hence, we abstract from such format switching.

[25] This treatment is further supported by the fact that the revenue share of the pharmacy segment is small relative to their CPG sales, a reverse of what we find at chain pharmacies.

[26] The formulation follows the potential-sites approach of Bodéré (2023), who nests a static spatial-demand model within a dynamic game over discrete locations. Unlike their setting, my model here has no endogenous quality investment. The binary occupancy structure is thus what preserves the finite-dependence collapse of Section 7.4. Moreover, I do not have to determine prices at the site level, since payers set dispensing prices and national firms set retail prices. Therefore, I do not worry about the direct computational burden of the neural-network approximation of the value function.

from start-of-year carryover occupancy, which is robust to same-year replacement overlap) to account for shopping malls. Capacity is one at the overwhelming majority of sites. I also include k_s, the median square footage of the site's occupants, so that vacant sites can carry a size that entrants can evaluate.

I assume that the three small-box types compete for the same pool of premises. A vacated pharmacy box can host an independent or a dollar store, and conversely. At the same time, big-box and grocery formats (B, G) occupy a disjoint large-box pool and evolve exogenously as described above. I also assume that the universe contains only ever-occupied locations, not the full set of developable parcels. I observe enough vacancy in the universe at any point in the sample period that I believe the assumption does not risk limiting entry by a pure lack of vacancy. Still, entry onto never-observed parcels is outside the model. This is innocuous for counterfactuals operating on the re-occupancy margin under moderate policy changes,[27] but could be binding under large ones that may induce large net entry.

The occupancy state. Let $\omega_{st} \in \{0, C, I, P\}$ denote the occupancy of site s in year t (a vector of slot occupancies where $q_s > 1$). The market state is the configuration together with the exogenous variables,

$$\omega_{mt} \;=\; \big(\omega_{st}\big)_{s \in \mathcal{S}_m}, \qquad x_{st} \;=\; x\big(s,\, \omega_{mt},\, z_{mt},\, m_{r,mt}\big), \tag{21}$$

where x_{st} collects the *local features* through which the configuration enters decisions: competitor counts by type within distance bands of the site, distance to the nearest competitor of each type, big-box and grocery presence, population within the same bands, county demographics z_{mt}, the reimbursement margin $m_{r,mt}$, site capacity and size, and year. The margin $m_{r,mt}$ is the dispensing margin per prescription, measured at the county-year level.[28]

Strategies are restricted to condition on x_{st} rather than on the full configuration ω_{mt}. This restriction is substantive, and I state it as the equilibrium concept: the solution is a moment-based Markov equilibrium in the sense of Ifrach and Weintraub (2017), in which firms best respond to the payoff-relevant summary of the configuration rather than to every element of it. With tens of thousands of sites, an exact Markov-perfect equilibrium on the configuration space is neither computable nor economically necessary given the atomistic role of any single distant store. Therefore, distance bands are chosen so that the features nest the county-level state, and the fit of the feature-based choice probabilities is reported. Firms hold rational expectations over the resulting first-order Markov transition of the state, which combines the exogenous evolution of

[27] In fact, the counterfactual rarely observes net entry.

[28] The county-year margin combines CMS Medicare and Medicaid fee schedules with county enrollment and is anchored to the national NCPA margin series $\bar{m}_{rt}$ (Appendix A.4). County-years without coverage carry the national value and contribute no cross-sectional variation. I exploit this decomposition of the margin into the national path and the county deviation to identify the dollar scale (Section 7.5).

demographics, the reimbursement margin, and the large-box pool, the reduced-form dollar-and-convenience process, and the equilibrium entry and exit behavior of the strategic types. As in Aguirregabiria and Suzuki (2014), each store treats the industry state as an exogenous Markov process when solving its own stopping problem.

Discussion: (near) uniform reimbursement. I treat reimbursement as (almost) uniform across pharmacies in a county-year and assume that the margin $m_{r,mt}$ is common to all pharmacies within the market. Doing so delegates state and federal-level policies—state Medicaid fee schedules, the federal DIR reform, county payer mix—rather than store-specific contract terms as the identifying variation for the supply parameters. This is a substantive assumption about a market in which reimbursement is negotiated between pharmacies and PBMs. In fact, its major threat is vertical integration, in which an integrated PBM reimburses rival pharmacies less generously than it does its own stores (Yde, 2025). I discuss why the caveat matters less here than it first appears in Appendix C.1. On the demand side, the assumption is innocuous at the relevant margin, again as the vast majority of prescriptions are filled with insurance and the consumer's out-of-pocket price does not vary across retailers (Appendix B.3).

Timing and the entry process. Within period t, (i) the exogenous states realize—demographics and the large-box configuration evolve, payers set $m_{r,mt}$; (ii) incumbents draw private i.i.d. extreme-value payoff shocks and choose either to stay or exit, with exit being terminal for the establishment and the scrap value normalized to zero (Section 7.3); (iii) each slot vacant at the start of t hosts *one* potential entrant, who draws an action-specific shock for each element of $\{\varnothing, C, I, P\}$ and either stays out or enters as one of the three formats, paying the type-specific entry cost $\kappa^{E,a}$. Slots that are vacated in t become available at $t+1$. The one-entrant-per-slot is the standard resolution of the coordination multiplicity that arises in simultaneous-entry games with entry choices smoothed by private information (Seim, 2006). It also maps directly into a multinomial conditional choice probability, which is how I estimate entry probabilities. I treat ownership transfers as non-events (the site persists; see Appendix A.3), format conversion as an exit followed by re-entry, and dollar-store entry as a reduced-form arrival process. The P branch of the entry probability is estimated and used in the simulation, but no structural entry cost is associated with it. Structural costs are therefore recovered for $\tau \in \{C, I\}$ only.[29]

[29] I assume that stores are the players: a chain store maximizes its own value rather than the chain's portfolio value. This is a strong assumption, especially in dense urban markets, where closing one chain store could divert demand to a nearby sibling store. The estimated exit probabilities then reflect the chain's internalization of cannibalization, in which a store-level objective is attributed to the site. Still, sites are substantially distant within a given chain, so their closures would likely divert consumers to a competing chain rather than a sibling store. Modeling chain pharmacy behavior as cannibalization within the duopoly (CVS and Walgreens) remains an aspiring future exercise.

7.3 Payoffs, Shocks, and the Scale of Unobservables

Variable profits by store type. The variable profit of store j combines retail margins and the dispensing margin,

$$\pi_j^* = \sum_{k\in\{x,f,p\}} \left(p_j^k - mc_j^k\right) Q_j^k \ + \ m_r\, Y_j, \tag{22}$$

where Y_j is the prescription volume. The dispensing term is zero for stores without pharmacies, and the retail term is zero for pharmacy-only independents. This lets the reimbursement policy shift profits directly to pharmacy operators, but not at all for those of retail-only formats. Fitting the retail component and fill volume separately, and reattaching the dispensing margin as m_r times the fitted volume, keeps the reimbursement channel explicit: a change in the margin moves the profit through fill volume alone. All dollar-denominated payoff objects are expressed in 2024 dollars using a single CPI-U deflator.

An active type-τ occupant of site s earns the per-period variable profit delivered by the demand model of Section 4: the sum of retail margins on the three retail categories and the dispensing margin on prescription volume,

$$\pi^{*,\tau}(x_{st}, k_j) \ = \ \pi^{\text{ret},\tau}(x_{st}, k_j) \ + \ m_{r,mt}\, Y^{\tau}(x_{st}, k_j), \tag{23}$$

where Y^{τ} is annual fills. Both components are known functions of the state given the demand estimates, and no supply-side parameter enters (23). The store-level profits underlying Section 7.3 are matched to sites and projected on the local features, separately by type, and separately for the retail component $\pi^{\text{ret},\tau}$ and the fill volume Y^{τ}, so that the payer margin enters as $m_{r,mt}\, Y^{\tau}(x_{st})$.[30]

Operating requires a fixed cost with market, size, and site components,

$$\phi^{\tau}(s) \ = \ \phi_0^{\tau}(m) \ + \ \phi_1^{\tau} k_s \ + \ \rho_s, \qquad \rho_s \ = \ g(r_s)'\gamma \ + \ \xi_s, \tag{24}$$

where $\phi_0^{\tau}(m)$ varies freely across markets, capturing rents and other local cost conditions, r_s collects observable site-cost shifters (local density, commercial rent proxies, distance to the nearest big-box store), and ξ_s is residual site quality. The coefficients γ on *observable* site-cost shifters are estimable. Still, site-specific structural costs are not—ξ_s is absorbed as an unobservable, in the spirit of a demand-side ξ, and the share of residual variance it accounts for is reported rather than interpreted. Also, ϕ^{τ} carries *no time variation* within a market. The mechanism attributes the closure wave to the collapse of π^*, and a flexible time path in ϕ^{τ} would absorb that variation. Entry requires a one-time cost $\kappa^{E,\tau}$ in addition to the first period's fixed

[30] Profit-store coordinates are block-group centroids while sites are establishment coordinates, so the match is the nearest same type store within a small radius, one-to-one up to the site's occupancy count. The match rates and distance diagnostics accompany the estimates. Importantly, the projection is evaluated at *vacant* sites as well, which prices the entry options for every format at every location, including the CPG-only profitability of a vacated pharmacy corner.

cost. The scrap value at exit is normalized, $\phi^{X,\tau} \equiv 0$, a normalization whose role in identification I return to below.

Each period, conditional on the state, the store draws action-specific private payoff shocks—continue and exit for incumbents; the format-specific entry options and staying out for potential entrants—that are i.i.d. type-I extreme value, and enter payoffs with scale $\sigma_\varepsilon^\tau > 0$ dollars. The firm is a risk-neutral dollar maximizer throughout, so σ_ε^τ should be interpreted not as a preference parameter but as the dollar dispersion of payoff components the firm observes that the econometrician does not.[31] Its presence is an empirical necessity rather than a convenience: a deterministic dollar rule implies degenerate exit rates at every observed state, whereas observed hazards move smoothly across states, with both exits and continuations at the same (x_{st}, k). The smoothness of the choice probabilities in dollar-denominated profit is what measures the scale. I work with the inverse of the scale,

$$\psi^\tau \;\equiv\; 1/\sigma_\varepsilon^\tau, \tag{25}$$

the sensitivity of choice indices to a one-dollar change in profit. Because $\pi^{*,\tau}$ arrives in dollars from the consumer panel through estimated demand, ψ^τ is an estimable parameter rather than a normalization.

7.4 Estimating Equations: Finite Dependence

The obstacle in any dynamic estimation is the continuation value. Computing it by forward simulation, in the tradition of two-step estimators of dynamic games (Bajari et al., 2007; Pakes et al., 2007; Aguirregabiria and Mira, 2007), requires the joint law of motion of the entire state and a stance on firms' beliefs about the demographic and reimbursement paths. In fact, this stance would continue under the reimbursement counterfactual; otherwise, it would require re-estimation. Worse, the sample period is a transition: imposing stationarity on observed market configurations would load the secular decline of π^* into the estimated costs, explaining exits with high fixed costs rather than with falling profits, which is, in fact, the hypothesis under test. Assuming finite dependence avoids the struggle with the continuation value (Arcidiacono and Miller, 2011; Almagro and Domínguez-Iino, 2025).

I exploit the finite-dependence properties of the action structure to deliver the collapse of the continuation value. Exit is terminal with normalized scrap value, and a site that exits returns to the entrant pool rather than disappearing. The option renewal matches the data construction in which the establishment record is the site and ownership transfers are not exit events (Appendix A.3). Applying the Hotz–Miller inversion

[31] This dispersion may include lease re-negotiations, owner-specific outside options, acquisition offers, local shocks absent from x_{st}, etc.

(Hotz and Miller, 1993) at the site level, the incumbent's stay/exit log-odds collapse to

$$\ell_\tau^X(x_{st}, k_j) \;=\; \psi^\tau\big[\pi^{*,\tau}(x_{st}, k_j) - \phi^\tau(s)\big] \;+\; \beta\,\mathbb{E}\big[-\ln P_\tau^X(x_{s,t+1})\,\big|\,x_{st}\big], \tag{26}$$

and the potential entrant's choice of format $a \in \{C, I\}$ against staying out collapses to

$$\ln\frac{P_a^E(x_{st})}{P_\varnothing^E(x_{st})} \;=\; \psi^a\big[\pi^{*,a}(x_{st}, k_s) - \phi^a(s) - \kappa^{E,a}\big] \;+\; \beta\,\mathbb{E}\big[\ln P_\varnothing^E(x_{s,t+1}) - \ln P_a^X(x_{s,t+1})\,\big|\,x_{st}\big], \tag{27}$$

where $\ell_\tau^X(x,k) \equiv \ln\big[(1 - P_\tau^X(x,k))/P_\tau^X(x,k)\big]$, P_τ^X is the type-τ exit probability, P_a^E the multinomial entry probability of action a at an open slot, and $P_\varnothing^E$ the stay-out probability. The continuation corrections are one-period-ahead functions of the *site-level* choice probabilities. The incumbent's correction concerns only its own next-period exit probability, and the entrant compares its post-entry exit probability with the option value of waiting in the same slot. The derivation of the inversion is detailed in Appendix C.2.

From (26)–(27), I build the continuation corrections from log-probabilities that are already in shock-scale units: they enter with β, with no ψ^τ attached. The discount factor is not identified from choice data (Magnac and Thesmar, 2002), and so I impose $\beta = 0.95$ rather than estimate a free coefficient on the continuation term. Then the system is linear in $(\psi^\tau, \psi^\tau\phi^\tau, \psi^\tau\kappa^{E,\tau}, \gamma)$ given first-stage choice probabilities: estimation requires no value function iteration, no equilibrium fixed point, and no forward simulation.

For estimation, I replace the conditional expectations with the fitted choice probabilities evaluated at the *realized* next-period state $x_{s,t+1}$—the realized-state construction of Kalouptsidi et al. (2021) applied at site granularity, with the exiting store's own occupancy restored in $x_{s,t+1}$ so that the correction is evaluated on the stay path. Under rational expectations, the expectational innovation, including rivals' responses, moves to the error term and is mean-independent of every time-t regressor.

7.5 Identification

The system (26)–(27) identifies its parameters in layers that differ in what they demand of the data, and I estimate them in that order.

Entry costs in index units. Differencing (27) from (26) *within a county-year* eliminates every common term: $\pi^{*,\tau}$ and the continuation corrections are functions of the cell-level state, and the market fixed cost $\phi_0^\tau(m)$ is common to the stay and entry margins of the same market. Hence, it differences out along with them. What survives is the scale-free entry cost

$$\kappa_u^\tau \;\equiv\; \psi^\tau\,\kappa^{E,\tau}, \tag{28}$$

estimated as (minus) the entrant contrast in a regression of the log-odds on an entrant indicator and size with market×year fixed effects and the site observables $g(r_s)$ as controls, so that the contrast compares entrants and incumbents at observably similar premises. Identification of κ_u^τ is thus immune to the measurement of π^*, to the level of fixed costs, and to the shock scale. No potential-entrant pool needs to be constructed by rule because the entry risk set is the observed count of physical open slots.

The dollar scale. ψ^τ is what connects index units to dollars, and it is identified only where dollar-denominated profits vary in the market. The cross-section alone cannot supply this variation. Under free entry, markets with high π^* also carry high rents $\phi_0^\tau(m)$. This places $\phi_0^\tau(m)$ in the error term for a pooled regression of ℓ^X on π^*, which is correlated with the regressor and is biased toward zero. I also cannot rely solely on aggregate time variation because the national margin path is confounded with aggregate hazard trends across the pandemic and policy reform years.

The design, therefore, isolates payer-driven, market-specific variation that the two-price separation of the demand side (Appendix B.3) makes a pure profit experiment, since no demand-side object responds to the reimbursement margin. While the drug margin still affects consumers indirectly through the front-store prices of the pricing game that internalizes fill traffic through the $m^r\,\partial Y_F/\partial p_F$ term of (20), it is closed at this level of variation. I assume chains price nationally, which implies that a state-year fee shock cannot affect the retail prices that a county's consumers face through the same institution that underlies the price indices, thereby providing a local profit shifter independent of demand. Moreover, the year effects absorb any response of national prices to the national margin path. I can then decompose predicted profit into a national-margin component and a payer deviation,

$$\pi_{st}^{*,\tau} \;=\; \underbrace{\pi_{st}^{\mathrm{ret},\tau} + \bar{m}_{rt}\,Y_{st}^\tau}_{\pi_{st}^{0,\tau}} \;+\; \underbrace{\big(m_{r,mt} - \bar{m}_{rt}\big)\,Y_{st}^\tau}_{\text{payer deviation}}, \tag{29}$$

where $\bar{m}_{rt}$ is the national margin. The deviation term is a shift-share object: state-year fee-schedule shocks (the shifts) scaled by the store's fill volume (the share). Substituting (29) into (26), imposing β, and replacing the conditional expectation of the continuation term with its realization $\hat{C}_{s,t+1}$ yields the estimating equation

$$\ell_{st}^X \,-\, \beta\,\hat{C}_{s,t+1} \;=\; \psi^\tau\,\pi_{st}^{0,\tau} \,+\, \psi^\tau\big(m_{r,mt}-\bar{m}_{rt}\big)Y_{st}^\tau \,-\, \psi^\tau\phi_1^\tau k_s \,-\, \psi^\tau g(r_s)'\gamma \,-\, \underbrace{\psi^\tau\phi_0^\tau(m)}_{\text{market FE}} \,+\, \delta_t \,+\, e_{s,t+1}. \tag{30}$$

Under rational expectations, the innovation in the realized continuation is mean-independent of every time-t regressor, including persistent components of $m_{r,mt}$, so it belongs in the error term (Kalouptsidi et al., 2021). The market fixed effects absorb $\phi_0^\tau(m)$ and all permanent profitability differences, and the year effects

δ_t absorb the aggregate hazard trend. The surviving variation is then the state-year fee changes common to a market and within-market differences in *exposure* to those changes because two stores in the same county-year face the same fee schedule but fill different volumes. Fee schedules are set at the state-payer level, so the exogeneity argument for the shocks is untouched by the finer unit. Several further layers of the argument—the role of the realized continuation, the recovery of fixed costs as by-products of (30), the re-occupancy margin that the site formulation adds, and the at-risk convention required by the censored site universe—are discussed in Appendix C.3.

Objects to recover for the counterfactual. The reimbursement counterfactual of Section 8 changes m_{rt} and nothing else, shifting every choice index by $\psi^\tau \, \Delta m \cdot Y^\tau(x,k)$. In the re-solved equilibrium, the cost terms ϕ^τ, $\kappa^{E,\tau}$, and the scrap normalization cancel from the CCP mapping. By the counterfactual-identification results of Aguirregabiria and Suzuki (2014), policies that shift only variable profit are invariant to the cost normalization, and their computation requires exactly ψ^τ, β, the baseline CCPs, the exogenous-state paths, and Y^τ. The cost estimates therefore serve the validation and interpretation purposes—the chain-versus-independent contrast in κ_u^τ—not the policy computation.

7.6 Estimation, Inference, and Aggregate Fit

Estimation proceeds in the order dictated by the layering above. The site panel is constructed from the establishment panel of Appendix A.3. I fit the first stages—type-specific exit logits, a multinomial entry logit over $\{\varnothing, C, I, P\}$, and the store-level profit projections evaluated at every site-year and type—on the local features x_{st}. The entry-cost contrasts κ_u^τ then come from the within-cell regression. Finally, I run the offset regression (30) to recover the scale, fixed costs, and site-costs. I provide further details for each step in Appendix C.4, including the requirement that the first-stage choice probabilities condition on the reimbursement margin.

First-stage fit. The cost estimates and the counterfactuals both run on the first-stage choice probabilities. Figures 18 and 19 assess the first-stage fit using dimensions the specifications do not target. Figure 18 groups observations into quantile bins of the fitted probability, and plots the observed frequency against the mean fitted probability within each bin for each exit logit and each entry alternative. The points lie on the forty-five-degree line across the probability range on all six margins. Regressing observed frequencies on bin means gives slopes between 1.03 and 1.07 for exit and between 1.03 and 1.17 for entry, and the fitted aggregate entry shares match the observed shares to within 0.005 percentage points.

Figure 19 examines the geography of exit, which is the dimension relevant for the access question. County identity enters the first stages through only a small set of county covariates, with no county effects, yet summing the fitted exit probabilities within each county reproduces the cross-county distribution of closures. The correlation between predicted and observed cumulative exits is 0.97 for chains, 1.00 for independents, and 0.99 for the dollar-and-convenience format. Because large counties accumulate both predicted and observed exits, the figure also reports the correlation of exit rates among counties with at least twenty site-years: 0.52, 0.34, and 0.67, respectively.

7.7 Cost Estimates

Table 10 reports the cost estimates.

Entry costs in index units. The within-county-year entrant contrast gives $\kappa_u^C = 8.36$ and $\kappa_u^I = 5.33$ (Table 10, Panel A). Perhaps unsurprisingly, chain entry costs exceed those of independents by a factor of roughly 1.6 in shock-scale units.

The scale. With flows calibrated by type, the payer-deviation coefficient is positive and statistically significant for both types: $\psi^C = 1.28$ and $\psi^I = 2.09$ per million dollars. The implied shock scales are plausible: $\sigma_\varepsilon^C = \$0.78$ million is of the order of per-store annual variable profit, and $\sigma_\varepsilon^I = \$0.48$ million scales sensibly with the smaller format.

Fixed costs: construction. The operating fixed cost is a by-product of the incumbent equation (26) rather than a free parameter given the scale estimate. Imposing $\hat{\psi}^\tau = \psi^\tau$ and rearranging, the operating-cost *index* $\phi_u^\tau \equiv \psi^\tau \phi^\tau$ is the mean Bellman residual over incumbent site-years,

$$\hat{\phi}_u^\tau \;=\; \mathbb{E}_w\Big[\hat{\psi}^\tau\,\hat{\pi}^{*,\tau}(x_{st},k_j) \;+\; \beta\gamma_E \;-\; \big(\hat{\ell}_\tau^X(x_{st},k_j) - \beta\,\hat{C}_{st}\big)\Big], \qquad \hat{\phi}^\tau \;=\; \hat{\phi}_u^\tau\,\hat{\sigma}_\varepsilon^\tau, \tag{31}$$

where $\hat{C}_{st} = -\ln \hat{P}_\tau^X(x_{s,t+1})$ is the realized continuation correction of Section 7.4, $\gamma_E \approx 0.577$ restores the Euler constant of the extreme-value shocks, and $\hat{\pi}^{*,\tau}$ is evaluated at the realized county reimbursement margin. No fixed effects are removed: ϕ^τ is a *level* by design, consistent with the identifying restriction that fixed costs carry no time variation (Section 7.3). Some caveats bound the interpretation: most notably, the level is identified only jointly with the scrap normalization $\phi^{X,\tau} = 0$ (Aguirregabiria and Suzuki, 2014).[32] It should be interpreted as mean site costs with residual site quality ξ_s absorbed. It also inherits the scale multiplicatively, both through $\hat{\psi}^\tau\hat{\pi}^*$ inside the residual and through $\hat{\sigma}_\varepsilon^\tau$ in the dollar conversion.

[32]Given that most closing pharmacies, especially independents, sell their list of patients to a nearby pharmacy, this normalization is likely a strong one, and would lead to an overestimation of the fixed costs. This is likely part of what explains the validation exercise placing the estimated fixed costs at the upper end of the audit for independents.

Fixed costs: estimates and takeaways. At these scales, the operating fixed costs are $\phi^C = \$3.78$ million and $\phi^I = \$1.25$ million per store-year, and the implied entry costs are $\kappa^{E,C} = \$6.5$ million and $\kappa^{E,I} = \$2.6$ million (Table 10, Panel B). As expected, chains are the heavier format on both margins—entry and operation each cost roughly three times the independent level in dollars—but the *ratio* of the two costs is stable across types: the entry barrier equals roughly 1.7 years of operating cost for chains and 2.0 for independents, so the cross-type difference in dynamics reflects overall cost scale rather than a differently shaped entry technology. The payer margin moves exit and entry in the direction the policy debate presumes, and the counterfactuals of Section 8 can use the dollar magnitudes.

External validation of the dollar scale. The cost levels can be compared with accounting data the estimation never touches, and the comparison doubles as a check on the scale itself, since the dollar value of the fixed cost is the index-unit Bellman residual converted at $\hat{\sigma}_\varepsilon^\tau$. For chains, Walgreens' financial statements imply a per-store operating cost—gross profit less net income, spread over roughly 8,600 U.S. stores—of \$3.5 million in fiscal 2023 and \$4.1 million in fiscal 2024, bracketing the estimated $\phi^C = \$3.78$ million. One-time impairment and litigation charges inflate the 2024 figure, so the two years plausibly bound the underlying operating cost. For independents, the NCPA Digest reports average per-store sales of \$5.4 million at an 18.2 percent gross margin, which at the industry's roughly two percent net margin implies about \$0.9 million of operating cost per store, and the income statements of the proprietary operator (Section 2) put direct per-store operating costs between \$0.35 million—a delivery-only satellite—and \$1.28 million, with most stores between \$0.70 and \$0.95 million. The estimated $\phi^I = \$1.25$ million lies within the observed range, toward its upper end.

The counterfactuals do not use the cost levels (Section 8.1). However, the cost estimates' agreement with independently observed accounting data is an overidentifying check on the dollar conversion ψ^τ that the counterfactuals do use.

8 Counterfactuals

I use the estimated model to evaluate the reimbursement-centered policy proposals that motivate the paper. The exercise fixes the *policy* and asks what it delivers: the per-prescription dispensing margin m_r is permanently restored to the highest level it has attained in real terms, and I compute what the policy costs, how many pharmacies it adds and where, and what the induced change in retail availability is worth to consumers in dollars. Throughout, I ask whether policies that operate on drug profitability can avert *desert-creating* closures at a cost commensurate with their welfare benefits, or whether they mostly transfer rents to stores in markets that are already well served.

8.1 Design: Stationary Equilibria and the Treatment of Beliefs

The counterfactual changes the margin process and nothing else, and it is computed on the site-level supply model of Section 7. A permanent margin change Δ shifts every pharmacy operator's flow profit by $\Delta \cdot Y^\tau$, in proportion to fill volume, and therefore shifts the stay and entry indices (26)–(27) by $\psi^\tau \Delta Y^\tau$ plus the induced change in the site-correct continuation corrections. I re-solve the entry/exit game for its long-run configuration under the shifted margin process by forward simulation at the estimated choice probabilities, with the policy and the induced competition both entering as structural index offsets scaled by ψ^τ, initialized at the data CCPs as the equilibrium-selection rule. I then compare long-run site configurations and their county aggregates, holding the exogenous large-box pool, demographics, retail profits $\pi^{\text{ret},\tau}$, and the demand primitives fixed.

By the counterfactual-identification results of Aguirregabiria and Suzuki (2014) and Kalouptsidi et al. (2021), computing this equilibrium requires exactly the objects the design identifies—ψ^τ, β, the baseline CCPs, the exogenous-state paths, and fill volumes Y^τ—and is invariant to the fixed-cost levels and the scrap normalization; the cost estimates of Section 7.7 serve as validation and interpretation, not computation. Chains and independents respond on both the exit and the entry margin; big-box, grocery, and CPG-only formats are unaffected by construction, since their presence is governed by their retail business, although big-box pharmacies still *collect* the higher margin on their fills, which matters for the fiscal accounting below.

Beliefs require a stance that the estimation deliberately avoided taking: the finite-dependence estimator recovers $(\psi^\tau, \kappa_u^\tau)$ without specifying what firms expect the reimbursement path to be, so there is no estimated belief process for the counterfactual to inherit. I impose the strongest available stance: the policy is *permanent, immediately credible, and unanticipated*, and I compare the stationary equilibrium it induces with the baseline stationary equilibrium. To the extent operators discount permanence, the count responses reported below are upper bounds on what a dollar of enacted policy achieves.

I make four additional assumptions that frame the results. I compare long-run configurations, abstracting from the transition path. I freeze the environment—demographics, the large-box pool, and retail profitability—at current fundamentals.[33] I fix the dollar scales that come from the shift-share design (Table 10), whose exposure attenuation runs conservative for the policy. Lastly, I hold copays and shelf prices fixed, so consumers are affected only through availability, and the welfare statements are transfer accounting with a consumer-surplus term, not a full incidence analysis of the drug price changes. Appendix C.5

[33] Again, the pharmacy segment is very small compared to the CPG segment for the large-box retailers. It remains a descriptive exercise to provide reduced-form evidence that large-box retailers compete among themselves and are only marginally influenced by small-box retailers' decisions.

elaborates on these limitations, especially on the near-equivalence of paid and credibly announced margin increases and why credibility cannot be assumed for free.

8.2 Counterfactual: Restoring the Margin to Its Real-Terms Peak

The national margin series $\bar{m}_{rt}$, expressed in 2024 dollars, attains its maximum at the start of the sample—2013 in the deflated NCPA series, at \$18.0 per fill against \$13.9 in 2024—and has eroded since. The policy sets

$$\Delta^{\text{peak}} \;\equiv\; \max_t \bar{m}^{\text{real}}_{rt} \;-\; \bar{m}^{\text{real}}_{r,2024}, \tag{32}$$

and adds Δ^{peak} to every county's margin permanently, preserving the county-level payer deviations. This is the strongest version of the proposals in current policy debates: no reform on the table restores more than the historical peak of drug profitability.[34]

I compare three configurations for the results. The first is the *observed* configuration—the pharmacy network as it stands in 2024. The second is the *baseline long run*: the configuration to which the estimated entry/exit process converges if the current margin process and current fundamentals persist. That is, the observed decline continued to its resting point. The third is the *peak-margin long run*, the resting point under Δ^{peak} but with the current fundamentals otherwise. The gap between the first two is what happens *without* the policy, and the policy effect is the gap between the last two. Comparing the counterfactual to the observed state would instead credit the policy with, or charge it for, the decline that occurs anyway during the transition to the baseline. Hence, every policy number below, including the welfare calculation, compares the two long-run configurations. Table 11 reports the outputs in this order. The count response and its location, the fiscal outlay, and the consumer value of the change in availability are discussed in the paragraphs below.

Store counts (Panel A). Exit falls and entry rises, both in proportion to fill volume, so the stationary pharmacy count weakly increases. The economic content lies in the magnitude and the location. Without the policy, the estimated process runs the 2024 network of 55,788 chain and independent pharmacies down to a long-run resting point of 29,723—a continued decline of 26,065 stores. Because chain exits are driven by the retail shortfall, which the margin does not touch, and because $\psi^{\tau}\Delta^{\text{peak}}Y^{\tau}$ is small relative to the dispersion of the payoff shocks for most stores, the count response to the policy is modest against that backdrop. The long-run national count rises by 9.3 percent relative to the no-policy baseline with 2,756 stores, of which

[34] Most of the proposals are at the state level, as Medicaid and even commercial insurance are operating at the state level. A counterfactual can be made more granular geographically to support a targeted subsidy argument. Still, I present the results for this strongest policy option primarily to highlight the welfare consequences and the no-policy baseline.

1,059 are chains and 1,696 independents, barely recovering a tenth of the continued decline. The induced stores concentrate where fills are plentiful, which is where pharmacies already operate: 98.1 percent of the induced store-mass accrues to counties with three or more pharmacies (as of 2024), while counties with at most one pharmacy, where the desert margin binds, gain 16 stores in total, and counties with two pharmacies gain 37.

Distance. Figure 20 plots the population-weighted distribution of distance to the nearest pharmacy under each configuration, similar to Figure 6, holding geography and population fixed. Hence, the curves move only through the pharmacy network. A decline that looks enormous in counts is barely visible in distances. Removing nearly half the network of 26,000 stores raises the population-mean distance from 1.8 to 2.1 miles, the median from 0.8 to 0.9 miles, and the share of the population living more than 10 miles from a pharmacy from 2.1 to 2.9 percent. The reason is similar to the geographic pattern documented in Section 3, where exits are concentrated in areas with high pharmacy density. Therefore, the marginal closure removes a store whose next-nearest neighbor is nearby, and the surviving network, including the constant big-box counters, absorbs the load with little added travel. This is the geographic form of the welfare result below: with distances nearly unchanged, the travel component of the consumer cost of the decline is necessarily small.

Fiscal cost (Panel B). The policy pays $\Delta^{\text{peak}} = \$4.17$ on *every* fill, infra-marginal ones included, and because fill volume is need-driven, the outlay is essentially an arithmetic object. Then, the current national fill volume, repriced by Δ^{peak}, is s roughly \$30.1 billion per year. Only 0.8 percent of it is paid on fills at stores whose presence the policy causes, and the implied price per induced store-year is \$10.9 million, which is several times the one-time entry cost of the store it induces (Table 10), paid every year.

Welfare (Panel C). The demand model delivers an exact money metric for the availability change because utility is quasilinear in the outside good. Conditional on the taste draws ν, per-occasion expected utility is the logsum over the basket choice set, and the compensating variation of moving consumer i's choice set from the baseline $\mathcal{C}_{it}^0$ to the counterfactual $\mathcal{C}_{it}^1$ is

$$\mathrm{CV}_{it}(\nu) \;=\; \frac{1}{\alpha_i}\left[\ln\sum_{c\in\mathcal{C}_{it}^1} e^{\,w_i(c)-\Gamma_i(c)} \;-\; \ln\sum_{c\in\mathcal{C}_{it}^0} e^{\,w_i(c)-\Gamma_i(c)}\right], \tag{33}$$

integrated over the taste draws and aggregated with population weights. Quasilinearity makes (33) exact: compensating and equivalent variations coincide because category demands carry no income effects. Dividing by α_i converts the utility gap into dollars, consumer by consumer, so the aggregate weights the change in availability according to who actually values it.

I recompute the retailer choices under the baseline and counterfactual configurations and re-solve the basket-choice problem, with a single solve on the union choice set delivering both logsums in (33). A consumer whose nearest retailers do not move contributes zero through the location channels by construction. The same solve prices all three configurations simultaneously, so Panel C also reports the CV for the *continued decline*. The consumer cost of moving from the observed network to the baseline long run, the welfare-side counterpart of Panel A's (2)−(1) row and the number the introduction pairs against the policy's price tag—\$20.2 billion per year, of which \$2.1 billion is added travel (the modest distance shifts of Figure 20), less than \$0.01 billion is lost local availability, and \$18.1 billion arrives through equilibrium re-pricing (below).

Note that each CV is expressed as a share of annual consumer spending on the modeled categories at the covered retailers because consumer-surplus levels are not identified (the logsum carries an arbitrary constant, so only differences have meaning). The continued decline costs consumers the equivalent of 1.85 percent of their annual spending, and the peak-margin policy restores only 0.03 percent of it. The consumer value of local pharmacy access falls steeply with local pharmacy density—the demand-side counterpart of stylized fact (iii), so the aggregate consumer value of the policy is small because the induced stores locate where the marginal pharmacy is worth the least. The policy's consumer value is \$0.32 billion per year against a fiscal outlay of \$30.1 billion, which translates to only one cent of consumer surplus per public dollar.

Equilibrium re-pricing. Closures affect the pricing competition, so prices are re-solved under each configuration. The implied price changes enter the CV exactly through the same union-solve device (every alternative duplicated at the counterfactual prices), and the decline rows in Table 11 report the resulting decomposition. I find that the dominant channel is not travel but prices. Survivors—above all, the distance-insensitive online channel, which absorbs the displaced demand—re-optimize upward when a class of competitors thins out, and a price increase of one to two percent applied to the full expenditure base dwarfs the travel channel, even though it is nearly invisible per basket. However, there are caveats attached to the price channel. One best-response step under strategic complementarity understates the full equilibrium response. Prices are uniform and national, as in the data, so local market power in the thin markets where counts fall most is outside the model's resolution. Also, the level of the step inherits the procedure's numerical residual relative to observed prices: the observed-configuration anchoring of that residual and the fully iterated version, in which each period's configuration is re-priced and fed back into the entry and exit indices, are left for the next draft.

Before I conclude, I return to the caveat that the paper has consistently assumed away: the CV excludes the health value of adherence. Under the deterministic-fill regime, no fill is ever lost, so (33) prices travel, assortment, retail prices, and one-stop convenience, but not the forgone prescriptions around closures. Where a policy changes the desert margin itself, (33) understates its value, and the last pharmacy in a market has unboundedly large CV under a mandatory fill. On that margin, the CV remains a lower bound—though this does not change the policy's arithmetic, because the induced stores overwhelmingly locate in markets far from the desert margin.

9 Conclusion

Recent pharmacy closures have prompted proposals to raise the margins pharmacies earn on dispensing, on the premise that restoring drug profitability would keep pharmacies open where access is at risk. Whether such policies can work depends on why pharmacies close and where entry and exit respond—questions that require linking consumer shopping behavior to the dynamics of market structure.

In this paper, I develop a multi-category demand model in which prescriptions anchor shopping trips, and consumers face shopping complementarity, purchasing CPGs to offset the extra trip cost at another retailer. I then embed the demand results into a dynamic model of pharmacy entry and exit, in which reimbursement serves as a policy lever.

The answer is unfavorable to a margin-based policy. Closures stem from the collapse of the front-store retail business that historically subsidized dispensing, not from drug margins, and they concentrate in markets dense with competing pharmacies. That concentration persists in the model's projections: under current conditions, the pharmacy network shrinks by nearly half in the long run, yet distance to the nearest pharmacy barely changes, and the consumer cost is modest, borne mainly by the prices of surviving stores rather than by lost access. Reimbursement does little to change this trajectory; for the same spatial reason, a margin is paid in proportion to fills, and fills are plentiful where pharmacies already operate. Restoring margins to their historical peak would cost \$30 billion per year, recover a tenth of the projected decline, place 98 percent of the induced stores in already-served counties, and return only about one cent of consumer surplus per public dollar.

In this analysis of the pharmacy market, I take reimbursement rates as given rather than modeling the contracts among payers, PBMs, and pharmacies that determine them. Avenues for future research include connecting pharmacy market structure to that vertical layer, and incorporating the health value of adherence, which the welfare calculations here deliberately exclude.

References

Aguirregabiria, V. and Mira, P. (2007). Sequential Estimation of Dynamic Discrete Games. *Econometrica*, 75(1):1–53. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0262.2007.00731.x.

Aguirregabiria, V. and Suzuki, J. (2014). Identification and Counterfactuals in Dynamic Models of Market Entry and Exit. *Quantitative Marketing and Economics*, 12(3):267–304.

Allcott, H., Diamond, R., Dubé, J.-P., Handbury, J., Rahkovsky, I., and Schnell, M. (2019). Food Deserts and the Causes of Nutritional Inequality. *The Quarterly Journal of Economics*, 134(4):1793–1844.

Almagro, M. and Domínguez-Iino, T. (2025). Location Sorting and Endogenous Amenities: Evidence From Amsterdam. *Econometrica*, 93(3):1031–1071. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA21394.

Amstislavski, P., Matthews, A., Sheffield, S., Maroko, A. R., and Weedon, J. (2012). Medication Deserts: Survey of Neighborhood Disparities in Availability of Prescription Medications. *International Journal of Health Geographics*, 11:48.

Arcidiacono, P., Ellickson, P. B., Mela, C. F., and Singleton, J. D. (2020). The Competitive Effects of Entry: Evidence from Supercenter Expansion. *American Economic Journal: Applied Economics*, 12(3):175–206.

Arcidiacono, P. and Miller, R. A. (2011). Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity. *Econometrica*, 79(6):1823–1867. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA7743.

Atkin, D., Faber, B., and Gonzalez-Navarro, M. (2018). Retail Globalization and Household Welfare: Evidence from Mexico. *Journal of Political Economy*, 126(1):1–73.

Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating Dynamic Models of Imperfect Competition. *Econometrica*, 75(5):1331–1370. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0262.2007.00796.x.

Battles, J. (2026). Do Pharmacies Matter?

Bell, D. R. and Lattin, J. M. (1998). Shopping Behavior and Consumer Preference for Store Price Format: Why "Large Basket" Shoppers Prefer EDLP. *Marketing Science*, 17(1):66–88.

Bodéré, P. (2023). Dynamic Spatial Competition in Early Education: an Equilibrium Analysis of the Preschool Market in Pennsylvania.

Brown, E. C. F. and McCuskey, E. Y. (2020). The Implications Of Rutledge v. PCMA For State Health Care Cost Regulation.

Cao, Y., Chevalier, J. A., Handbury, J., Parsley, H., and Williams, K. R. (2024). Willingness to Travel with Endogenous Distance: Evidence from the Changing Retail Landscape.

Caputi, T. (2026). Patients Adapt: The Mitigated Impact of Retail Pharmacy Closures.

Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The Effect of Minimum Wages on Low-Wage Jobs. *The Quarterly Journal of Economics*, 134(3):1405–1454.

Cubanski, J. and Damico, A. (2023). Medicare Part D in 2024: A First Look at Prescription Drug Plan Availability, Premiums, and Cost Sharing.

CVS Health (2024). Independent Pharmacies: Myths versus Reality.

CVS Health Corporation (2024). CVS Health Corporation Reports Fourth Quarter and Full-Year 2024 Results.

Deaton, A. and Muellbauer, J. (1980). An Almost Ideal Demand System. *The American Economic Review*, 70(3):312–326.

DellaVigna, S. and Gentzkow, M. (2019). Uniform Pricing in U.S. Retail Chains. *The Quarterly Journal of Economics*, 134(4):2011–2084.

Deshpande, M. and Li, Y. (2019). Who Is Screened Out? Application Costs and the Targeting of Disability Programs. *American Economic Journal: Economic Policy*, 11(4):213–248.

Dubé, J.-P., Joo, J., and Kim, K. (2025). Discrete/Continuous Choice Models and Representative Consumer Theory. *Journal of Economic Theory*, 226:106012.

Federal Trade Commission (2024). Pharmacy Benefit Managers: The Powerful Middlemen Inflating Drug Costs and Squeezing Main Street Pharmacies. Technical report.

Fein, A. (2024). Economic Report on Pharmaceutical Wholesalers and Specialty Distributors.

Feldman, R. (2023). Pharmacy Benefit Managers and the Prescription Drug Supply Chain: Impact on Patients and Taxpayers. Technical report, Senate Finance Committee.

Gentzkow, M. (2007). Valuing New Goods in a Model with Complementarity: Online Newspapers. *American Economic Review*, 97(3):713–744.

Guadamuz, J. S., Alexander, G. C., Zenk, S. N., Kanter, G. P., Wilder, J. R., and Qato, D. M. (2021a). Access To Pharmacies and Pharmacy Services in New York City, Los Angeles, Chicago, and Houston, 2015-2020. *Journal of the American Pharmacists Association*, 61(6):e32–e41.

Guadamuz, J. S., Wilder, J. R., Mouslim, M. C., Zenk, S. N., Alexander, G. C., and Qato, D. M. (2021b). Fewer Pharmacies In Black And Hispanic/Latino Neighborhoods Compared With White Or Diverse Neighborhoods, 2007–15. *Health Affairs*, 40(5):802–811.

Hausman, J., Leonard, G., and Zona, J. D. (1994). Competitive Analysis with Differenciated Products. *Annales d'Économie et de Statistique*, (34):159–180.

Hess, J. D. and Gerstner, E. (1987). Loss Leader Pricing and Rain Check Policy. *Marketing Science*, 6(4):358–374.

Holmes, T. J. (2011). The Diffusion of Wal-Mart and Economies of Density. *Econometrica*, 79(1):253–302. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA7699.

Hortacsu, A. and Syverson, C. (2015). The Ongoing Evolution of US Retail: A Format Tug-of-War. *Journal of Economic Perspectives*, 29(4):89–112.

Hotz, V. J. and Miller, R. A. (1993). Conditional Choice Probabilities and the Estimation of Dynamic Models. *The Review of Economic Studies*, 60(3):497–529.

Houde, J.-F. (2012). Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline. *American Economic Review*, 102(5):2147–2182.

Ifrach, B. and Weintraub, G. Y. (2017). A Framework for Dynamic Oligopoly in Concentrated Industries. *The Review of Economic Studies*, 84(3):1106–1150.

IQVIA (2024). The Global Use of Medicines 2024.

Jia, P. (2008). What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry. *Econometrica*, 76(6):1263–1316. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA6649.

Kalouptsidi, M., Scott, P. T., and Souza-Rodrigues, E. (2021). Linear IV Regression Estimators for Structural Dynamic Discrete Choice Models. *Journal of Econometrics*, 222(1, Part C):778–804.

Leung, J. and Li, Z. (2022). Big-Box Store Expansion and Consumer Welfare.

Magnac, T. and Thesmar, D. (2002). Identifying Dynamic Discrete Decision Processes. *Econometrica*, 70(2):801–816. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-0262.00306.

McFadden, D. (1989). A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration. *Econometrica*, 57(5):995–1026.

Mulligan, C. B. (2023). Ending Pay for PBM Performance: Consequences for Prescription Drug Prices, Utilization, and Government Spending.

NASHP (2023). State Laws Passed to Lower Prescription Drug Costs: 2017–2024.

NCPA (2025). 2025 NCPA Digest: State of the Industry.

Nowosielski, B. (2024). Independent Pharmacies Continue to Face Financial Hardships as the Clock Ticks on PBM Reform. 02.

Pakes, A., Ostrovsky, M., and Berry, S. (2007). Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples). *The RAND Journal of Economics*, 38(2):373–399. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1756-2171.2007.tb00073.x.

Pakes, A. and Pollard, D. (1989). Simulation and the Asymptotics of Optimization Estimators. *Econometrica*, 57(5):1027–1057.

Qato, D. M., Alexander, G. C., Chakraborty, A., Guadamuz, J. S., and Jackson, J. W. (2019). Association Between Pharmacy Closures and Adherence to Cardiovascular Medications Among Older US Adults. *JAMA Network Open*, 2(4):e192606.

Qato, D. M., Daviglus, M. L., Wilder, J., Lee, T., Qato, D., and Lambert, B. (2014). 'Pharmacy Deserts' Are Prevalent In Chicago's Predominantly Minority Communities, Raising Medication Access Concerns. *Health Affairs*, 33(11):1958–1965.

Qato, D. M., Zenk, S., Wilder, J., Harrington, R., Gaskin, D., and Alexander, G. C. (2017). The Availability of Pharmacies in the United States: 2007–2015. *PLoS ONE*, 12(8):e0183172.

Rhodes, A. (2015). Multiproduct Retailing. *The Review of Economic Studies*, 82(1):360–390.

Rite Aid Corporation (2024). Rite Aid Corporation Reports Fiscal 2024 First Quarter Results.

Schneier, C. (2025). *Distributional Effects of Exclusive Dealing in Retail Real Estate*. PhD thesis, University of Chicago.

Seim, K. (2006). An Empirical Model of Firm Entry with Endogenous Product-Type Choices. *The RAND Journal of Economics*, 37(3):619–640. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1756-2171.2006.tb00034.x.

Smith, H. (2004). Supermarket Choice and Supermarket Competition in Market Equilibrium. *The Review of Economic Studies*, 71(1):235–263.

Song, I. and Chintagunta, P. K. (2007). A Discrete–Continuous Model for Multicategory Purchase Behavior of Households. *Journal of Marketing Research*, 44(4):595–612.

Starc, A. and Swanson, A. (2021). Preferred Pharmacy Networks and Drug Costs. *American Economic Journal: Economic Policy*, 13(3):406–446.

Sweeting, A. (2013). Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry. *Econometrica*, 81(5):1763–1803. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA7473.

Thomassen, O., Smith, H., Seiler, S., and Schiraldi, P. (2017). Multi-category Competition and Market Power: A Model of Supermarket Pricing. *American Economic Review*, 107(8):2308–2351.

Wales, T. J. and Woodland, A. D. (1983). Estimation of Consumer Demand Systems with Binding Non-Negativity Constraints. *Journal of Econometrics*, 21(3):263–285.

Walgreens Boots Alliance (2024). Walgreens Boots Alliance Reports Fiscal 2024 Third Quarter Results.

Wittenauer, R., Shah, P. D., Bacci, J. L., and Stergachis, A. (2022). Pharmacy Deserts and COVID-19 Risk at the Census Tract Level in the State of Washington. *Vaccine: X*, 12:100227.

Wittenauer, R., Shah, P. D., Bacci, J. L., and Stergachis, A. (2024). Locations and Characteristics of Pharmacy Deserts in the United States: A Geospatial Study. *Health Affairs Scholar*, 2(4):qxae035.

Yde, E. (2025). Vertical Integration and Regulated Profits in Pharmacy Benefits Markets.

Ying, X., Kahn, P., and Mathis, W. S. (2022). Pharmacy Deserts: More Than Where Pharmacies Are. *Journal of the American Pharmacists Association: JAPhA*, 62(6):1875–1879.

Figures

Figure 1: Number of Retail Pharmacies by Type, 2017–2024

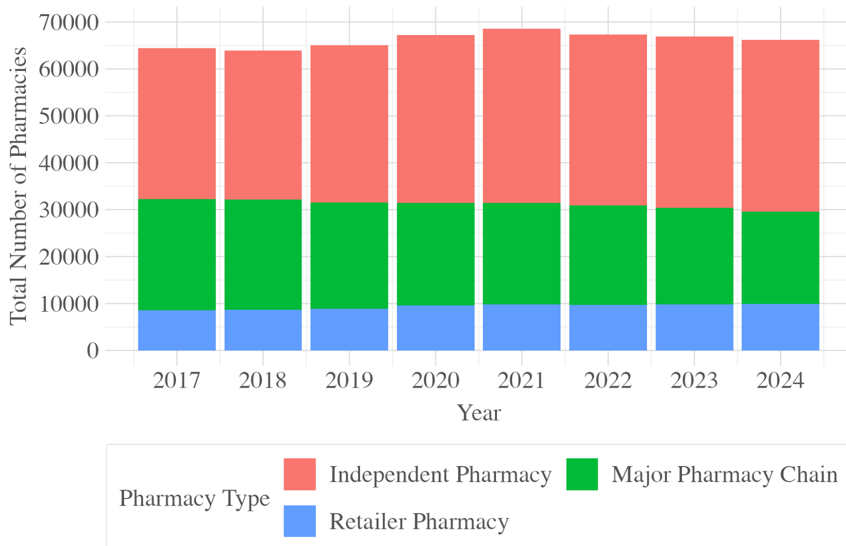

Notes: Counts are from the establishment panel of Appendix A.3, net of churning (simultaneous openings and closings), with ownership transfers bridged so that rebrandings and divestitures are not counted as closures. Major pharmacy chains are CVS, Walgreens, and Rite Aid. Retailer pharmacies are big-box and grocery stores that operate pharmacy counters. Independent pharmacies are the remaining retail, outpatient-facing pharmacies.

Figure 2: Number of Retail Pharmacies by Type, Indexed to 2017

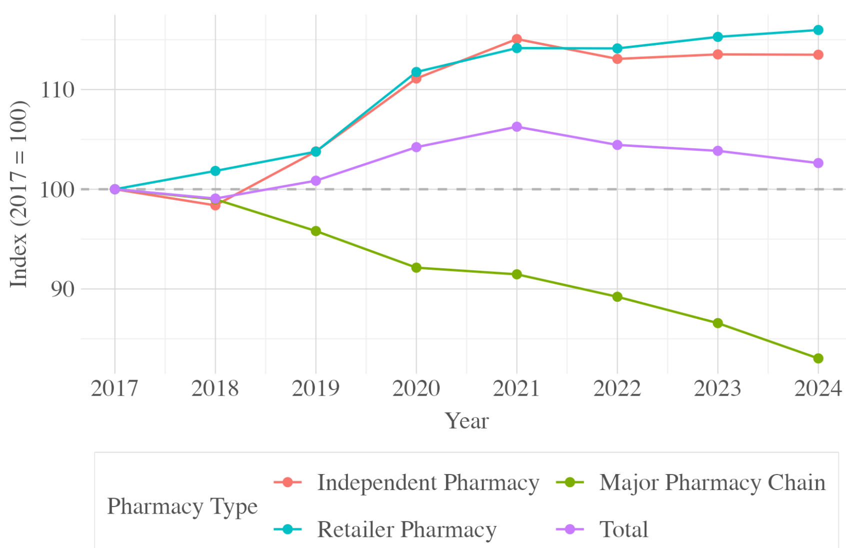

Notes: Each series divides the type's pharmacy count by its 2017 value. Counts are constructed as in Figure 1.

Figure 3: Effect of Pharmacy Closure on Monthly Days of Prescription Supply

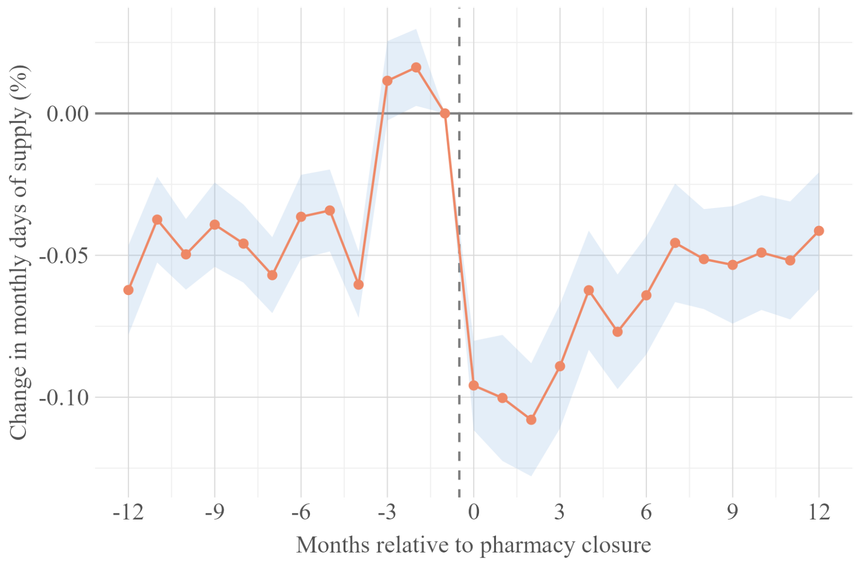

Notes: Stacked Poisson event-study estimates of equation (1) using MarketScan commercial claims. Coefficients are proportional effects on monthly days of prescription supply relative to the month before closure ($k = -1$, omitted). Treatment is the closure of a pharmacy–insurance identifier that serves at least thirty distinct patients and is followed by at least twelve months of zero claims. A patient is assigned to a closing pharmacy if more than half of her fills in the three months before closure occur at that pharmacy. Controls are patients whose own pharmacy closes at least twelve months after the event window, and the sample is restricted to patients continuously enrolled with drug coverage over the event window. The outcome counts all of a patient's fills at any pharmacy, including those with a missing pharmacy identifier, and through mail order. Bands are 95% confidence intervals, with standard errors clustered on the pharmacy identifier.

Figure 4: Net Change in Pharmacies by Initial County Pharmacy Count, 2021–2024

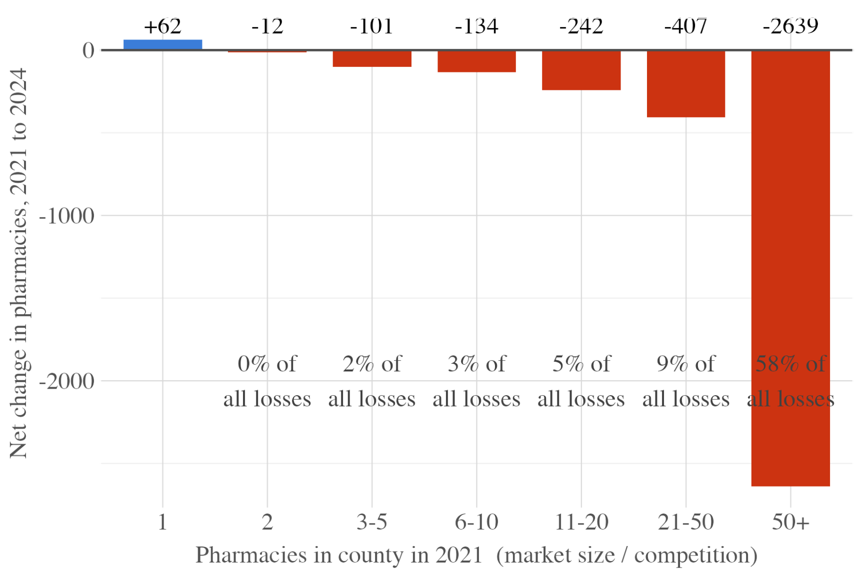

Notes: Counties are binned by their 2021 pharmacy count. Labels above the bars report the net change in the number of pharmacies in each bin between 2021 and 2024. Labels within the panel report each bin's share of all net losses. Source: establishment panel of Appendix A.3.

Figure 5: County-Level Net Change in Pharmacies versus the 2021 Count

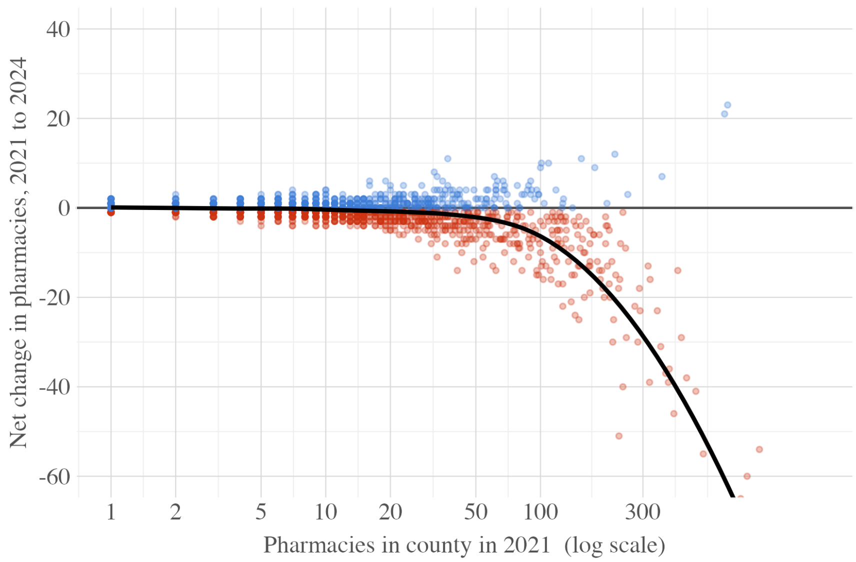

Notes: Each dot is a county. The horizontal axis is the county's 2021 pharmacy count on a log scale. The vertical axis is the county's net change in pharmacies between 2021 and 2024. The line is a flexible fit. Source: establishment panel of Appendix A.3.

Figure 6: Distance to the Nearest Retail Pharmacy, 2021 versus 2024

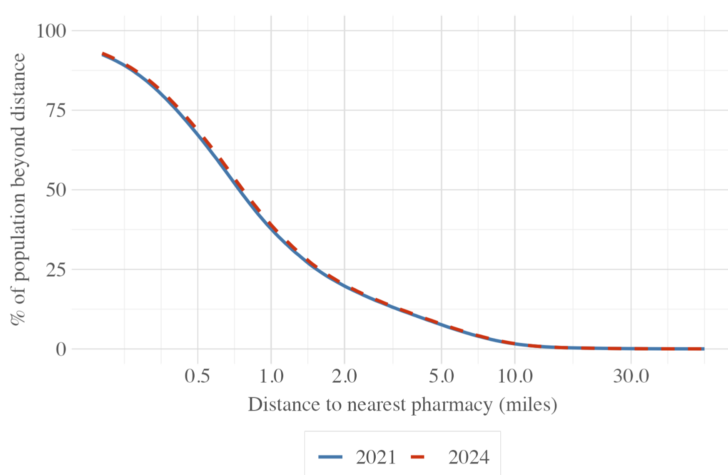

Notes: Share of the population living beyond each distance from the nearest retail pharmacy, in 2021 and 2024. Distance is measured from each census block group's centroid to the nearest retail pharmacy in the geocoded establishment panel of Appendix A.3, and the distribution weights block groups by population.

Figure 7: Driving Time to the Nearest Retail Pharmacy, 2021 versus 2024

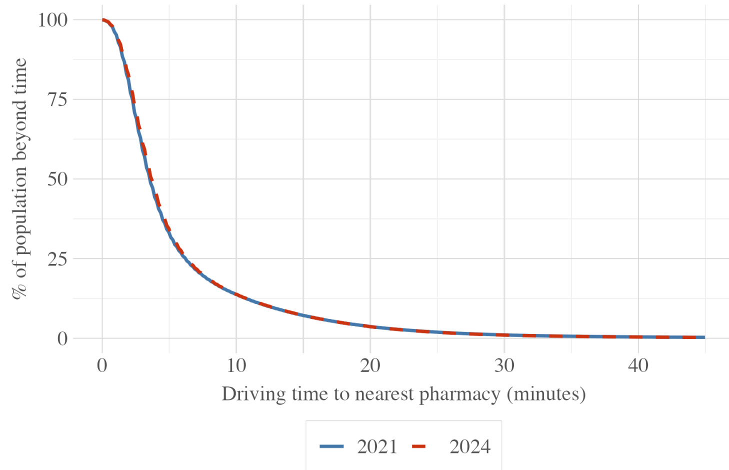

Notes: Share of the population living beyond each driving time from the nearest retail pharmacy, in 2021 and 2024. Construction mirrors Figure 6—block-group centroids to the nearest pharmacy, weighted by block-group population—with driving time in place of distance.

Figure 8: Average Pharmacy Gross Margin per Prescription, 2013–2024

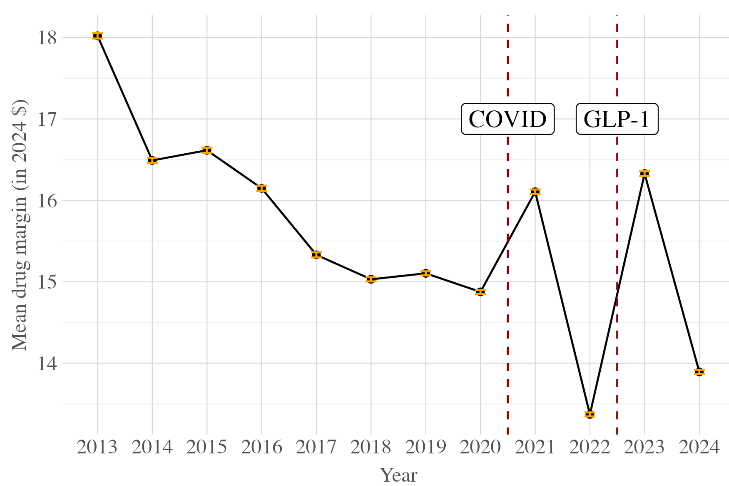

Notes: Volume-weighted mean gross margin per prescription: allowed amounts from MarketScan commercial claims net of acquisition costs from the NADAC series, deflated to 2024 dollars. Dashed vertical lines mark the onset of COVID-19 and the diffusion of GLP-1 drugs.

Figure 9: Pharmacy Gross Margin as a Percentage of the Prescription Price, 2013–2024

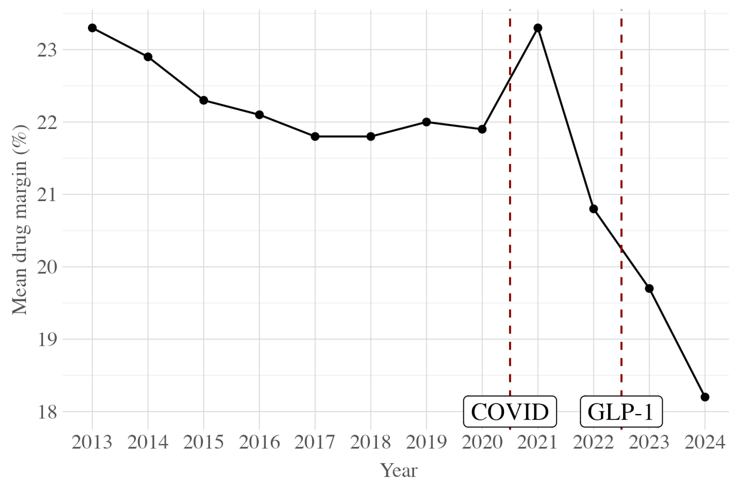

Notes: Volume-weighted mean gross margin per prescription, expressed as a percentage of the prescription price. Margins are constructed as in Figure 8.

Figure 10: Average Out-of-Pocket Payment per Prescription, 2013–2024

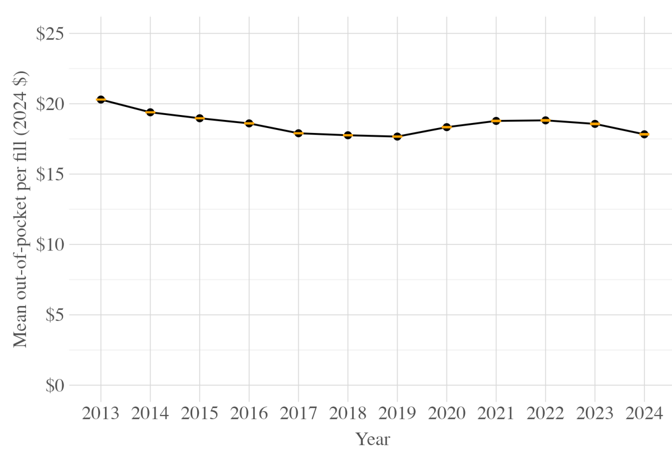

Notes: Mean out-of-pocket amount (copay, coinsurance, and deductible payments) per fill, MarketScan commercial claims, deflated to 2024 dollars.

Figure 11: Total CPG Expenditure at Chain Pharmacies

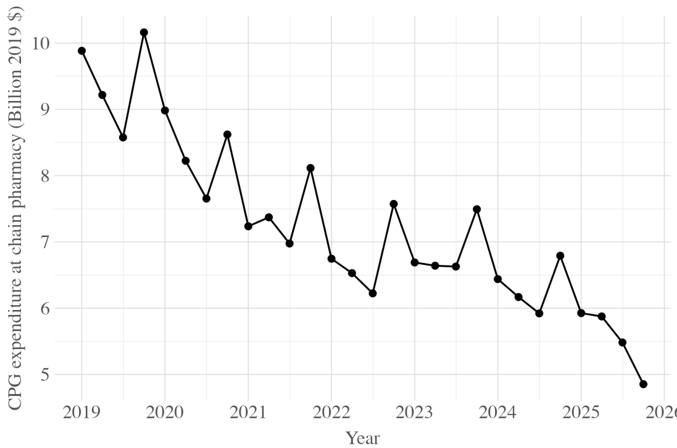

Notes: Quarterly total consumer spending at chain pharmacies (CVS, Walgreens, Rite Aid) accounted for by CPG purchases, projection-weighted, from the Numerator household panel.

Figure 12: CPG Expenditure Share at Chain Pharmacies

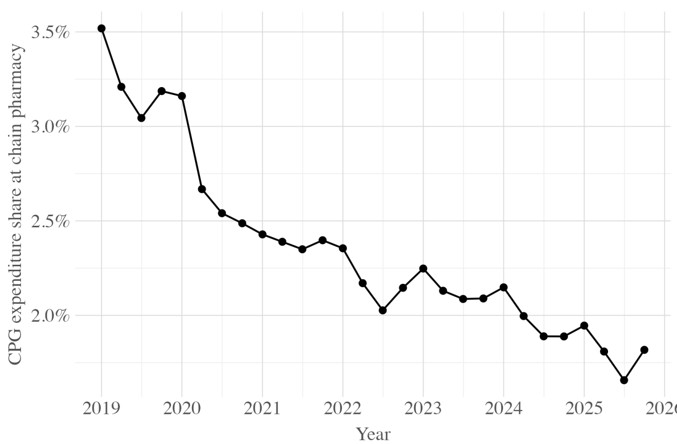

Notes: Quarterly share of consumer spending at chain pharmacies (CVS, Walgreens, Rite Aid) accounted for by CPG purchases, projection-weighted, from the Numerator household panel.

Figure 13: Share of Drug Trips at Chain Pharmacies Including a CPG Purchase

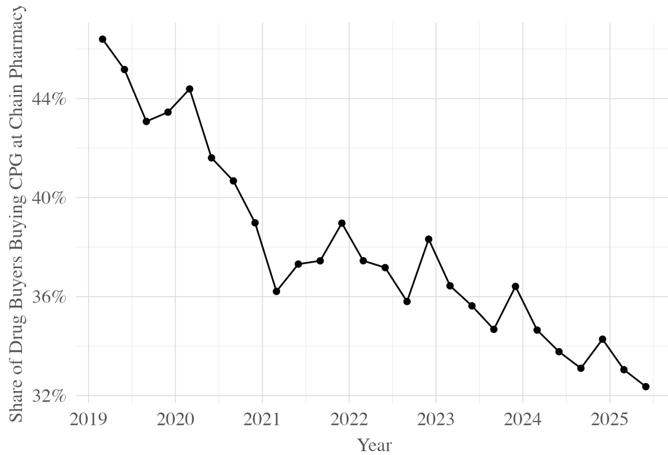

Notes: Among shopping trips to chain pharmacies that include a prescription purchase, the quarterly share that also includes at least one CPG item, from the Numerator household panel.

Figure 14: Pharmacy Type Composition by Market Size

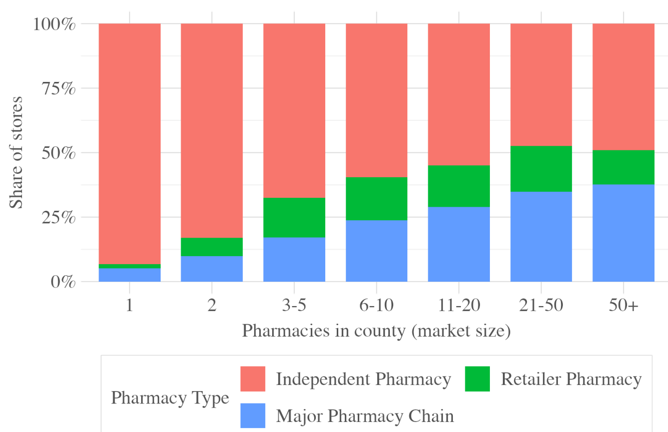

Notes: Counts are from the establishment panel of Appendix A.3, net of churning (simultaneous openings and closings), with ownership transfers bridged so that rebrandings and divestitures are not counted as closures. Major pharmacy chains are CVS, Walgreens, and Rite Aid. Retailer pharmacies are big-box and grocery stores that operate pharmacy counters. Independent pharmacies are the remaining retail, outpatient-facing pharmacies.

Figure 15: Decomposition of the Decline in Chain Front-Store Revenue, 2021–2024

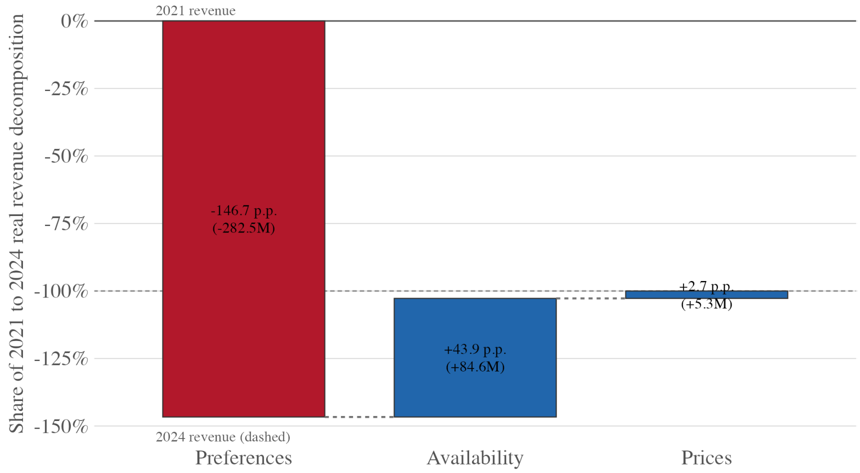

Notes: Sequential decomposition of the 2021–2024 decline in chain pharmacies' front-store revenue (the retail categories only; prescription fills and the dispensing margin are excluded), in constant 2024 dollars. Starting from the 2021 market, each bar replaces one channel—preferences, availability (the configuration of stores available to each household), and retail prices—with its 2024 value, holding the previously switched channels at 2024 and the remaining channels at 2021, and re-simulates front-store revenue. The dashed line marks the 2024 level, that is, 100 percent of the decline. The underlying demand solves are inflation-free: in mixed-year cells the price coefficient is rescaled by the CPI ratio, so the price bar is a relative-price effect rather than the mechanical passage of inflation. Bar labels report each channel's share of the total decline (percentage points) and the dollar change (millions, sample-weighted). Channels are cumulated in the fixed order shown—preferences, then availability, then prices—so each bar is that channel's marginal contribution conditional on the previously switched channels.

Figure 16: Distribution of Recovered Marginal Costs by Category, 2019 versus 2024

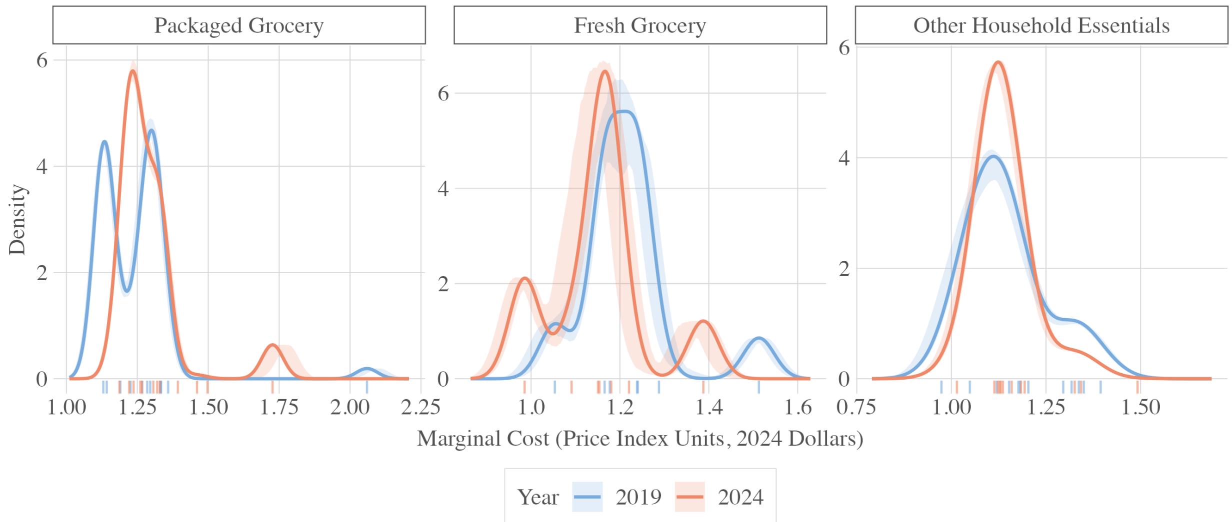

Notes: Kernel densities of the marginal costs recovered from the Bertrand–Nash first-order conditions (20), across banner–category products, weighted by predicted quantity. Prices are the banner–category hedonic indices in Appendix A.2, so costs are in index units, deflated to constant 2024 dollars using the CPI-U. The 2019 and 2024 densities are therefore comparable in real terms. Shaded bands are pointwise 90% confidence bands from a parametric bootstrap over the estimated demand-parameter distribution, holding the kernel bandwidth and the weights fixed at their point-estimate values. Tick marks on the horizontal axis locate the product-level point estimates. Chain pharmacies carry no fresh groceries, and the dollar-store and online formats do not operate in every category.

Figure 17: Distribution of Profit Margins by Category, 2019 versus 2024

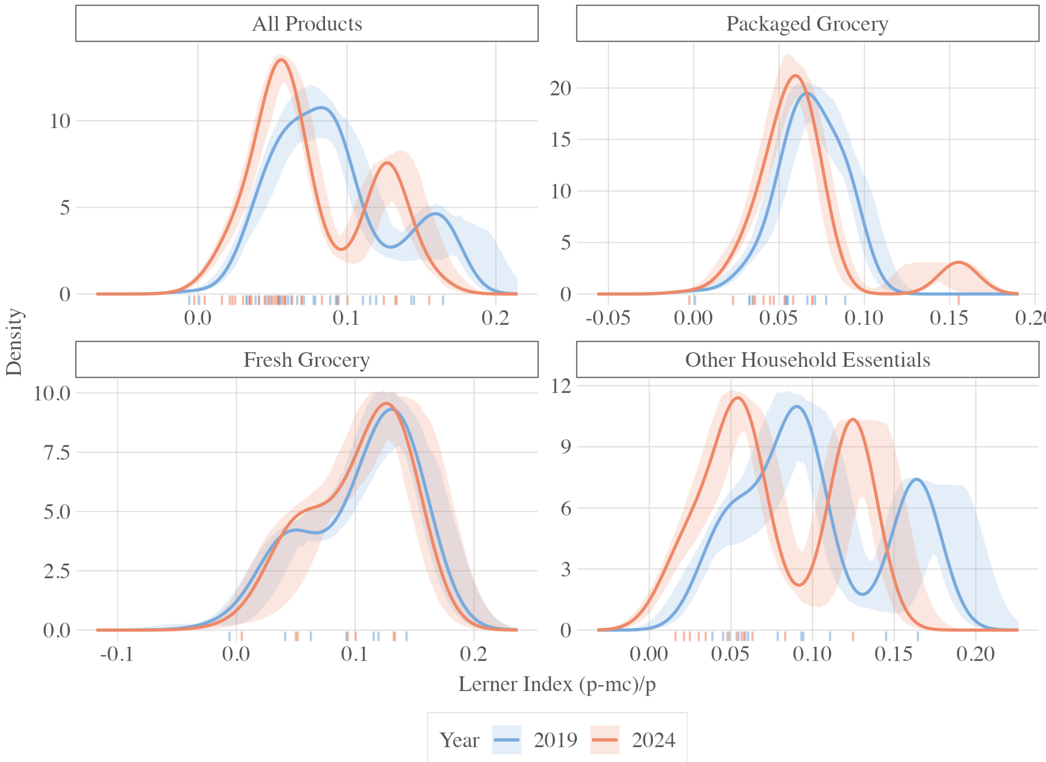

Notes: Kernel densities of the Lerner index $(p-mc)/p$ implied by the recovered marginal costs, across banner–category products, weighted by predicted revenue. The "All Products" panel pools the three categories. The Lerner index is a ratio, so it is invariant to the CPI deflation applied in Figure 16. Shaded bands are pointwise 90% confidence bands from a parametric bootstrap over the estimated demand-parameter distribution, holding the kernel bandwidth and the weights fixed at their point-estimate values. Tick marks on the horizontal axis locate the product-level point estimates. Format coverage is as in Figure 16.

Figure 18: Calibration of the First-Stage Choice Probabilities

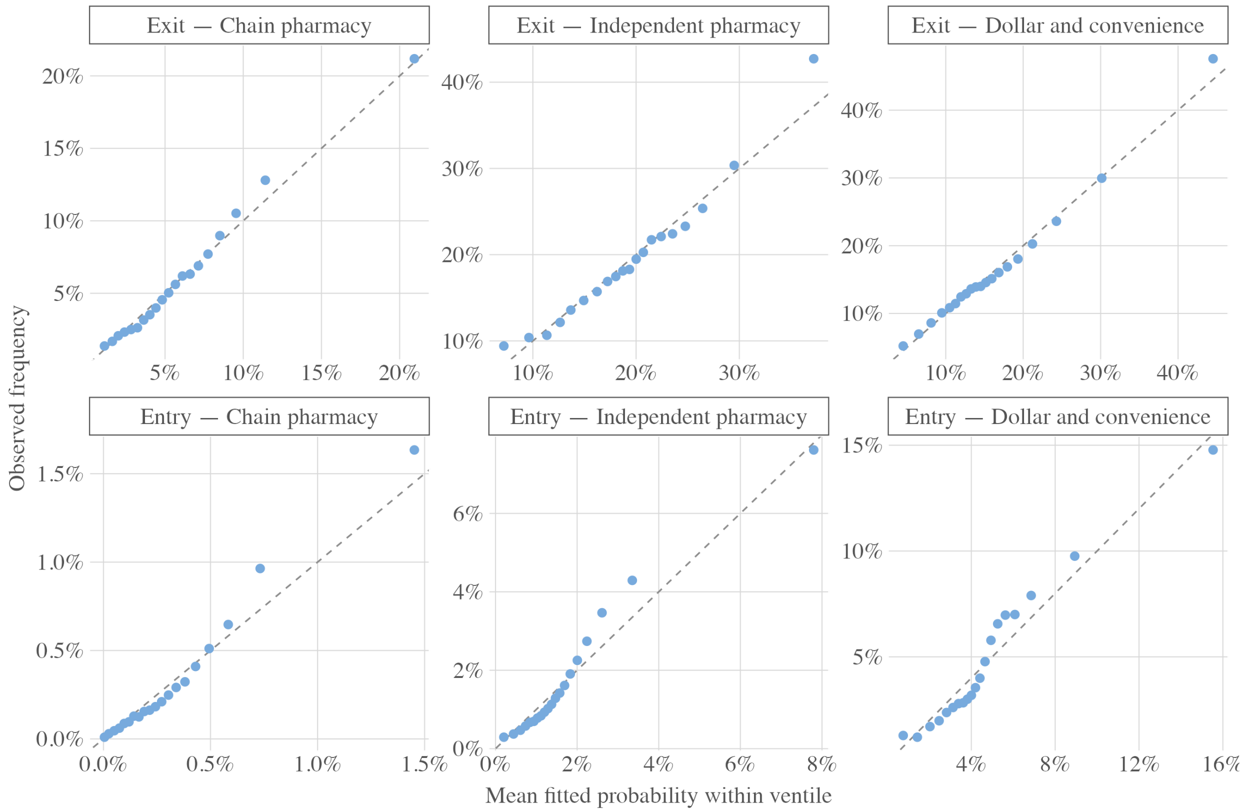

Notes: Each panel groups observations into twenty quantile bins of the fitted choice probability and plots the observed frequency within the bin against the mean fitted probability. The dashed line is the 45-degree line. Top row: the type-specific exit logits, fit on incumbent site-years. Bottom row: the three entry alternatives of the slot-level multinomial logit, fit on at-risk slot-years. Entry probabilities are an order of magnitude smaller than exit probabilities, and the dollar-and-convenience branch is the reduced-form arrival process of Section 7.2. Regressing observed frequencies on bin means yields slopes ranging from 1.03 (exit, dollar, and convenience) to 1.17 (entry, chain). The specifications do not target quantile bins.

Figure 19: The Geography of Exit: Predicted versus Observed Closures by County

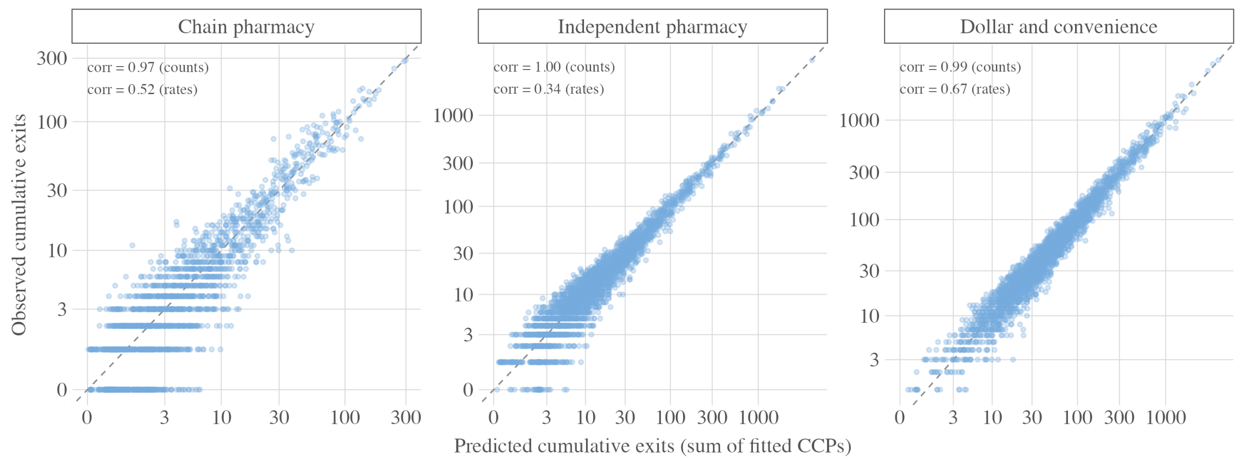

Notes: Each point is a county. The horizontal axis is the sum of fitted exit probabilities across the county's incumbent site-years. The vertical axis is the number of observed exits. Both axes are on a $\log(1 + x)$ scale, and the dashed line is the 45-degree line. County identity enters the exit logits only through county-level covariates (store counts by type and ACS demographics), not through county fixed effects, so the cross-county distribution of closures is not targeted. The count correlations partly reflect county size, since large counties accumulate both predicted and observed exits. The rate correlations compare predicted and observed exit hazards across counties with at least 20 site-years.

Figure 20: Distance to the Nearest Pharmacy under Counterfactual Configurations

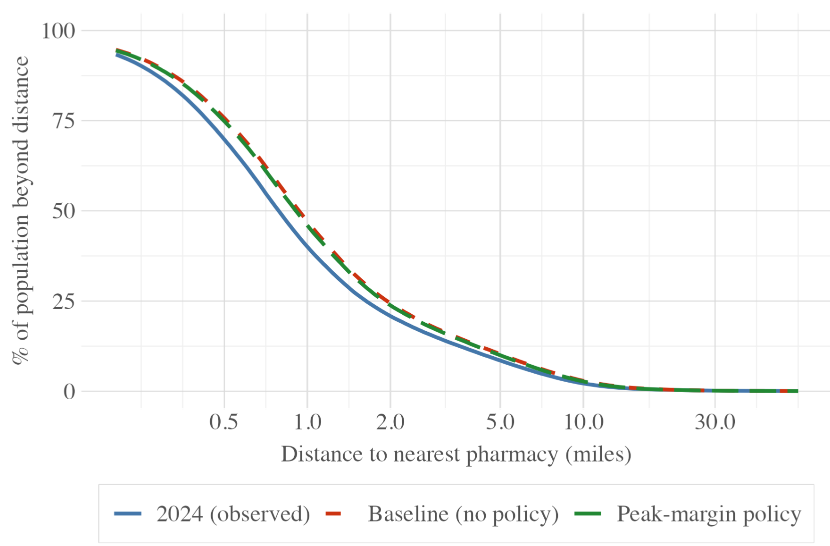

Notes: Population-weighted survival curves over 2020-vintage census block groups: the share of the population living beyond each distance from the nearest pharmacy under the observed 2024 network, the baseline long run (the decline run to completion), and the peak-margin long run. Geography and population (2024 weights) are held fixed, and big-box pharmacy counters—constant across configurations—are included, so the curves differ only through the chain and independent configuration.

Tables

Table 1: Net Pharmacy Availability and Prescription Consumption at the Market Level

	Prescription days supply (per enrollee-year)
Net Δlog pharmacy stock	0.003
(elasticity)	(0.018)
Effect of 10% net contraction	-0.03% of days
95% CI (elasticity)	$[-0.033,\ 0.039]$
Patient fixed effects	Yes
Year fixed effects	Yes
Clustering	MSA

Notes: Fixed-effects Poisson estimate of patient-year prescription days supplied on the net annual change in the log count of active pharmacies in the patient's MSA, MarketScan commercial claims merged to the establishment panel of Appendix A.3. The coefficient is the elasticity of prescription consumption with respect to net local pharmacy availability. The Poisson specification retains enrolled patients with zero fills. Standard errors, in parentheses, are clustered by MSA.

Table 2: Reduced-Form Drivers of Net Pharmacy Losses

	Net loss in pharmacies
Expenditure share	-0.131***
	(0.037)
Dual purchase share	0.003*
	(0.002)
Drug margin	0.072
	(0.164)
County fixed effects	Yes
Year fixed effects	Yes
Clustering	State

Notes: Fixed-effects Poisson estimates of equation (2) on the county-year panel, 2019–2024, with the county's lagged pharmacy count as the offset, so that coefficients are semi-elasticities of the net-loss rate. Regressors are lagged one year: the chain-pharmacy share of county CPG expenditure and the dual-purchase share (both in percentage points, from the household panel), and the county dispensing margin as a percentage of the average prescription charge. Standard errors, in parentheses, are clustered by state. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Estimation Sample: Demographics, Choice Environment, and Price Indices

Panel A. Demographic variables (3,000 consumers)	Mean	SD
Household size	2.66	1.21
Household income ($1,000s per year)	99.1	56.6
Household weekly income ($) per head	806.6	510.0

Panel B1. Choice environment: availability and nearest store within 30 miles (70,430 occasions)

		Available		Distance (miles)		Size (1,000 sq ft)		Price
Firm	Type of firm	≤15 mi	≤30 mi	Mean	SD	Mean	SD	Mean
CVS	Chain drugstore	0.92	0.96	3.4	4.6	12.6	10.5	1.18
Walgreens	Chain drugstore	0.94	0.98	3.5	4.6	15.9	11.6	1.31
Rite Aid	Chain drugstore	0.40	0.43	4.8	5.8	14.9	4.5	1.33
Independent drugstores	Drugstore	0.79	0.92	7.7	6.5	19.9	31.3	–
Walmart	Mass merchant	0.97	1.00	4.1	3.8	143.0	34.3	1.10
Target	Mass merchant	0.84	0.92	5.5	5.8	147.5	18.9	1.15
Costco	Club	0.65	0.77	8.2	6.8	145.5	39.0	1.26
Sam's Club	Club	0.67	0.82	9.2	6.9	143.4	38.8	1.16
Kroger	Food-and-drug	0.27	0.36	9.2	8.8	74.4	53.6	1.10
Dollar/convenience	Discount	0.99	1.00	1.9	1.9	10.7	9.1	1.17
Amazon	Online	–	–	–	–	–	–	1.65
Regional food-and-drug	Food-and-drug	0.84	0.93	5.2	6.3	91.7	68.3	1.25
Regional grocery	Grocery	0.87	0.93	4.5	5.1	34.9	43.8	1.15
Others	Small chains	0.99	1.00	2.0	2.7	12.0	26.4	1.08

Panel B2. Price indices by category (available store-week cells)

	Packaged	Fresh	Other	Drugs
Mean	1.28	1.19	1.14	1.00
Standard deviation	0.46	0.19	0.19	0.00
Standard deviation (within banner)	0.39	0.10	0.16	0.00

Notes: The sample is 3,000 households from the Numerator panel observed on randomly sampled days over 2019–2024 (70,404 shopping occasions; Section 5). Panel A reports household demographics. Panel B1 describes the choice environment: for each retailer, the share of occasions on which it is available within 15 and 30 miles, the distance to and floor space of the nearest outlet, and the mean price index across categories. Independent pharmacies carry no retail price index because they dispense prescriptions only, and the drug price index is one at all pharmacies by construction (Appendix B.3). Panel B2 summarizes the banner-week category price indices of Appendix A.2.

Table 4: Estimation Sample Construction

	Consumers	Occasions	Store visits	Expenditure share
0. Categorized purchases on sampled consumer-days	3,000	70,430	98,095	1.000
1. One drugstore kept on occasions with >2 drugstore banners	3,000	70,430	98,095	1.000
2. Top-two retailer cap (fill occasions keep all drugstores)	3,000	70,430	91,558	0.974
3. Sparse drugstore-category cells removed	3,000	70,430	91,389	0.972
4. One retailer per category (top by spending)	3,000	70,430	84,365	0.907
5. Estimation sample (demographics present)	3,000	70,430	84,365	0.907

Notes: Each row applies one screening step to the categorized purchase records on sampled consumer-days. Columns report the surviving consumers, occasions, store visits, and share of expenditure retained. The steps impose the basket structure of Section 4—at most two retailers per occasion and one retailer per category, with fill occasions retaining all pharmacies—and together preserve over 90 percent of expenditure.

Table 5: Shopping Outcomes in the Estimation Sample

Panel A1. Use of stores per shopping occasion (share of occasions, pre-screening)			
One store			0.72
Two stores			0.21
Three stores			0.05
Four or more stores			0.02
Occasion includes a prescription fill			0.04
Panel A2. Occasions with >1 store (3,980 occasions)		Mean	SD
Expenditure share in 1st store by daily spending (store A)		0.70	0.16
Expenditure share in 2nd store by daily spending (store B)		0.26	0.13
Panel A3. Share of category spending in 2nd store for the category, from (A, B)		Mean	SD
All occasions		0.03	0.08
Occasions with >1 store		0.12	0.11

Panel B. Category spending per occasion in the estimation sample

		Spending ($/occasion)		Zeros
Category	Illustrative products	Mean	SD	(share)
Packaged	Canned/frozen food, snacks, beverages, dairy	21.63	33.27	0.28
Fresh	Produce, meat, seafood, deli, bakery	10.82	20.31	0.48
Other	Personal care, OTC health, household	29.45	68.13	0.30
All categories		62.89	84.93	0.00

Panel C. Two-store occasions pairing a drugstore with another retailer (pre-screening)

		Drugstore main for category (share)		
Store pair (A, B)	Occasions	Packaged	Fresh	Other
Drugstore and grocery	267	0.08	0.01	0.73
Drugstore and dollar/online/other	148	0.47	0.21	0.42
Drugstore and mass merchant/club	139	0.17	0.05	0.29
Observations: 70,430 shopping occasions (3,000 consumers, 2019–2024)				

Notes: Panel A reports store use per shopping occasion before screening. Panel B reports category spending per occasion and the share of occasions with zero purchases in each category in the estimation sample. Panel C reports two-store occasions pairing a pharmacy with another retailer and the share of those occasions on which the pharmacy is the main store for each category. 70,404 occasions from 3,000 consumers, 2019–2024.

Table 6: Demand Estimates by Year

Year		2019		2020		2021		2022		2023		2024	
		Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE
Panel A. Store-category taste effects ($\boldsymbol{\beta}$) *and scaling terms* ($\boldsymbol{\sigma}$)													
Fresh Grocery	β_{01}	1.495	0.293	1.554	0.303	1.527	0.304	1.564	0.213	1.447	0.236	1.517	0.454
Packaged Grocery	β_{02}	0.501	0.060	0.483	0.048	0.466	0.049	0.495	0.020	0.498	0.069	0.501	0.060
Household size	β_1	0.132	0.025	0.197	0.022	0.159	0.026	0.172	0.012	0.178	0.016	0.169	0.028
ln(floor space)	β_2	0.005	0.001	0.007	0.001	0.006	0.001	0.007	0.001	0.006	0.001	0.005	0.001
Fixed across category/store	σ_1	0.411	0.027	0.424	0.024	0.420	0.058	0.423	0.065	0.409	0.036	0.432	0.101
Time-varying	σ_2	1.517	0.053	1.690	0.055	1.557	0.064	1.533	0.069	1.532	0.091	1.508	0.039
Category specific	σ_3	2.376	0.074	2.570	0.089	2.554	0.120	2.481	0.111	2.480	0.158	2.418	0.039
Store/category specific	σ_4	0.045	0.004	0.046	0.004	0.046	0.004	0.048	0.003	0.043	0.003	0.043	0.004
Panel B. Second-order quadratic parameters ($\boldsymbol{\Lambda}$)													
Fresh Grocery	λ_{ff}	0.247	0.045	0.246	0.047	0.250	0.048	0.251	0.037	0.249	0.042	0.238	0.068
Packaged Grocery	λ_{pp}	0.079	0.009	0.074	0.007	0.071	0.007	0.073	0.004	0.075	0.009	0.071	0.008
Other	λ_{xx}	0.087	0.003	0.088	0.003	0.082	0.004	0.085	0.004	0.081	0.005	0.077	0.002
Fresh-Packaged Grocery	λ_{fp}	0.011	0.001	0.016	0.002	0.020	0.002	0.019	0.002	0.001	0.002	0.009	0.002
Other-Fresh Grocery	λ_{xf}	0.003	0.001	0.001	0.001	-0.005	0.001	-0.005	0.002	0.003	0.001	0.005	0.002
Other-Packaged Grocery	λ_{xp}	0.002	0.001	-0.004	0.001	-0.003	0.001	-0.004	0.001	0.001	0.001	0.002	0.001
Fresh Grocery-Drug	λ_{fy}	0.056	0.308	0.052	0.344	0.052	0.261	0.052	0.320	0.054	0.288	0.058	0.277
Packaged Grocery-Drug	λ_{py}	0.452	0.136	0.422	0.121	0.412	0.117	0.409	0.115	0.413	0.121	0.422	0.113
Other-Drug	λ_{xy}	0.534	0.194	0.528	0.211	0.515	0.174	0.524	0.191	0.521	0.055	0.488	0.077
Panel C. Price parameters ($\boldsymbol{\alpha}$)													
Constant	α_1	0.318	0.028	0.335	0.036	0.364	0.035	0.329	0.024	0.342	0.012	0.314	0.022
Constant (2024 dollars)	$\alpha_1 \times \mathrm{CPI}_t/\mathrm{CPI}_{2024}$	0.259	0.023	0.276	0.029	0.314	0.030	0.307	0.023	0.332	0.011	0.314	0.022
1/[weekly income ($) per head]	α_2	0.341	0.039	0.349	0.038	0.357	0.073	0.345	0.042	0.358	0.025	0.338	0.040
Panel D. Shopping costs ($\boldsymbol{\gamma}$)													
Two store dummy	γ_{11}	4.715	0.150	5.110	0.231	4.742	0.119	5.097	0.217	4.595	0.112	4.466	0.102
Standard deviation	γ_{12}	0.100	2.309	0.100	3.892	0.100	1.866	0.100	3.200	0.100	1.817	0.100	1.729
Distance	γ_{21}	0.137	0.006	0.153	0.009	0.143	0.007	0.142	0.007	0.143	0.007	0.170	0.008
Standard deviation	γ_{22}	0.059	0.006	0.057	0.012	0.058	0.008	0.059	0.008	0.059	0.008	0.055	0.013

Notes: Simulated-method-of-moments estimates of the full random-coefficients model of Sections 4–5 under the by-year protocol: all parameters are re-estimated separately on each year's occasions, with winsorization caps and demographic centering constants frozen at their pooled 2019–2024 values so that units are comparable across columns. All specifications additionally include store-format-by-category fixed effects (ξ_{bk}), year-quarter time effects (ψ_t), and the pharmacy-format drug intercepts ($\xi_{b,y}$); these are estimated but not reported, as in Thomassen et al. (2017). The category scale is anchored by $\beta_{0x} \equiv 1$, and the basket-level extreme-value shock has unit scale, so utility parameters are expressed in units of the essentials category. The row "Constant (2024 dollars)" rescales α_1 by $\mathrm{CPI}_t/\mathrm{CPI}_{2024}$: the price indices are nominal (Appendix A.2), so α_1 is the marginal utility of a nominal year-t dollar. The converted row expresses it per constant 2024 dollar. The conversion touches no other parameter—quantities are deflated expenditures, so the remaining estimates carry quantity, distance, or utility units, and α_2 is invariant once incomes are likewise expressed in 2024 dollars. Standard errors are clustered by consumer and include the $(1 + 1/R)$ simulation correction of Section 5.

Table 7: Model Fit: Store-Category Demands

	In-sample		Out-of-sample	
Year	Quantity	Shoppers	Quantity	Shoppers
Panel A. Correlation between predicted and observed demands, $\rho(Q_{fk}, Q^*_{fk})$, $\rho(D_{fk}, D^*_{fk})$				
2019	0.986	0.990	0.972	0.985
2020	0.989	0.982	0.985	0.983
2021	0.977	0.971	0.974	0.970
2022	0.985	0.982	0.978	0.975
2023	0.986	0.979	0.986	0.976
2024	0.988	0.979	0.988	0.982
Panel B1. Firm share of category demand, $\lvert s_{fk} - s^*_{fk} \rvert$				
2019	0.007	0.007	0.012	0.008
2020	0.008	0.008	0.010	0.007
2021	0.009	0.009	0.010	0.009
2022	0.008	0.007	0.010	0.009
2023	0.009	0.008	0.009	0.009
2024	0.008	0.008	0.009	0.007
Panel B2. One-stop share of category demand, $\lvert s_{1ss,k} - s^*_{1ss,k} \rvert$				
2019	0.058	0.017	0.050	0.017
2020	0.076	0.022	0.074	0.022
2021	0.067	0.020	0.075	0.026
2022	0.071	0.022	0.061	0.020
2023	0.060	0.018	0.043	0.022
2024	0.052	0.019	0.055	0.021

Notes: Panel A reports the correlation between predicted and observed store-category quantity demands and shopper counts, by year, in the estimation sample and out of sample. Panel B1 reports the mean absolute error in firm shares of category demand; Panel B2 reports the mean absolute error in the one-stop share of category demand. Predictions are computed from the estimated model at the by-year parameters (Table 6). The out-of-sample columns evaluate the model on a holdout of 3,000 additional households drawn from the panel and processed identically to the estimation sample.

Table 8: Model Fit: Firm Revenue and Shopper Shares

	In-sample				Out-of-sample			
	Revenue		Shoppers		Revenue		Shoppers	
Firm	Pred	Obs	Pred	Obs	Pred	Obs	Pred	Obs
CVS	0.011	0.009	0.029	0.027	0.012	0.008	0.030	0.027
Walgreens	0.014	0.010	0.027	0.027	0.014	0.011	0.028	0.026
Rite Aid	0.002	0.002	0.006	0.005	0.002	0.002	0.006	0.005
Independent Drugstore	0.000	0.000	0.007	0.003	0.000	0.000	0.007	0.003
Walmart	0.214	0.200	0.162	0.166	0.219	0.209	0.167	0.177
Sam's Club	0.023	0.030	0.023	0.018	0.024	0.032	0.025	0.019
Costco	0.053	0.069	0.039	0.032	0.054	0.074	0.041	0.033
Target	0.066	0.066	0.057	0.057	0.068	0.068	0.059	0.060
Kroger	0.042	0.039	0.039	0.037	0.040	0.042	0.037	0.038
Regional Drug Grocery	0.136	0.132	0.128	0.118	0.135	0.123	0.128	0.114
Regional Grocery	0.050	0.061	0.082	0.063	0.051	0.059	0.085	0.063
Amazon	0.148	0.172	0.143	0.164	0.150	0.175	0.149	0.171
Dollar Store	0.022	0.016	0.039	0.050	0.007	0.009	0.012	0.030
Other Food Store	0.169	0.148	0.152	0.160	0.170	0.140	0.155	0.158
Other	0.051	0.047	0.069	0.073	0.053	0.049	0.072	0.075

Notes: Predicted and observed shares of revenue and of shoppers by firm, pooled over 2019–2024, in the estimation sample and out of sample. Firms are the banner-level alternatives of the demand system. "Regional Drug Grocery," "Regional Grocery," "Dollar Store," "Other Food Store," and "Other" are composite banners.

Table 9: Model Fit: Composition of Fill Locations

	In-sample		Out-of-sample	
Fill location	Pred	Obs	Pred	Obs
CVS	0.202	0.288	0.198	0.278
Walgreens	0.145	0.242	0.148	0.220
Rite Aid	0.009	0.059	0.007	0.057
Independent Drugstore	0.019	0.078	0.019	0.081
Walmart	0.207	0.077	0.210	0.089
Sam's Club	0.009	0.010	0.010	0.010
Costco	0.034	0.023	0.034	0.022
Target	0.051	0.054	0.052	0.050
Kroger	0.066	0.056	0.063	0.071
Regional Drug Grocery	0.258	0.113	0.259	0.122

Notes: Predicted and observed shares of prescription fills by pharmacy banner among fill occasions, pooled over 2019–2024, in the estimation sample and out of sample.

Table 10: Site-Level Supply Cost Estimates

	Chains (C)	Independents (I)
Panel A: estimates in shock-scale units		
Entry cost, κ_u^τ	8.36	5.33
	(0.04)	(0.03)
Dollar scale, ψ^τ (per \$1M)	1.28	2.09
	(0.28)	(0.34)
Panel B: implied dollar values (real 2024 \$M)		
Payoff-shock scale, $\sigma_\varepsilon^\tau = 1/\psi_{\text{dev}}^\tau$	0.78	0.48
Entry cost, $\kappa^{E,\tau} = \kappa_u^\tau \, \sigma_\varepsilon^\tau$	6.54	2.55
Fixed operating cost, ϕ^τ (per store-year)	3.78	1.25

Notes: Point estimates with analytic standard errors in parentheses: county-clustered for κ_u^τ, state-clustered for ψ^τ (the identifying fee shocks are state-level). Block-bootstrap standard errors propagating first-stage estimation error—including for the dollar values in Panel B—are not yet computed. $\beta = 0.95$ is imposed, and the scrap value is normalized to zero, so cost levels are net of annuitized scrap. Store flows are calibrated to type-specific external fill anchors (CVS/Walgreens filings for C, NCPA per-pharmacy prescriptions for I).

Table 11: The Peak-Margin Counterfactual

Panel A: pharmacy counts (chains + independents)	
(1) Observed count, 2024	55,788
(2) Long-run count, current policy (baseline)	29,723
(3) Long-run count, peak margin	32,478
Big-box pharmacies (type B; exogenous, constant)	37,134
Continued decline without the policy, (2)−(1)	−26,065 (−46.7%)
Policy effect, (3)−(2)	+2,756 (+9.3%)
of which chains / independents	+1,059 / +1,696
Induced stores in counties with ≤ 1 pharmacy	+16 (0.6% of total)
Induced stores in counties with 2 pharmacies	+37 (1.3%)
Induced stores in counties with ≥ 3 pharmacies	+2,703 (98.1%)
Panel B: fiscal cost (real 2024 dollars, per year)	
Margin increase Δ^{peak}, per fill (peak: 2013)	\$4.17
Outlay, $\Delta^{\text{peak}} \times$ equilibrium fills	\$30.1B
of which at big-box pharmacies	\$5.3B
Share paid on policy-induced fills	0.8%
Outlay per induced store-year	\$10.9M
Panel C: consumer welfare (per year)	
CV of the continued decline, (2) vs. (1)	−\$20.23B
of which: travel distance	−\$2.11B
of which: local availability	−\$0.00B
of which: equilibrium re-pricing	−\$18.12B
memo: added miles per affected fill occasion	+0.33
as a share of annual spending at the modeled retailers	−1.85%
CV of the policy, (3) vs. (2)	+\$0.32B
as a share of annual spending at the modeled retailers	+0.03%
Consumer surplus per public dollar	\$0.011

Notes: Three configurations: the observed 2024 network, the baseline long run—the resting point of the estimated entry/exit process under the *current* margin process and current fundamentals (i.e., the observed decline run to completion), and the peak-margin long run (the model of Section 7, fundamentals frozen at 2024). The policy effect is the gap between the two long-run configurations, so it is not credited with the decline that occurs in the absence of the policy. Induced-store rows classify counties by their 2024 pharmacy count. Big-box pharmacies are exogenous to the policy and constant across configurations; they remain available substitutes in the Panel C welfare calculation. The Panel B outlay reprices every fill at the higher margin. Because fill volume is need-driven, it is essentially the current national fill volume times Δ^{peak}, and nearly all of it is inframarginal. Panel C converts the availability change into dollars through the compensating variation (33), evaluated by re-solving the demand model on the baseline and counterfactual configurations (each consumer's nearest pharmacy outlet by banner recomputed under each), aggregated with population weights and anchored by occasion type: fill occasions to the external fill series the supply model uses, retail-only occasions to a national shopping-trips benchmark—the panel's undercoverage is specific to drug trips, and a single fill-based factor would overstate the retail side by an order of magnitude. Percentage rows normalize the CVs by annual consumer expenditure on the modeled categories at the covered retailers, using the same anchors. Consumer-surplus *levels* are not identified in a quasilinear model (only differences are), so shares of expenditure are the interpretable scale. Retail prices are re-optimized for each configuration using a single damped best-response step in the national pricing game, with the recovered marginal costs accounted for. The decline rows decompose the CV into travel-distance, availability, and re-pricing channels, and the re-pricing—crude by construction (one step, uniform national prices)—inherits the step's numerical residual against observed prices, with the anchored and fully iterated versions deferred. The count and welfare responses are conservative under exposure attenuation of ψ^{τ} (Table 10), and the compensating variation excludes the adherence-health margin, on which it remains a lower bound.

A Data Appendix

This appendix describes the construction of the datasets the paper builds from raw sources: the estimation categories in the household panel, the store-category price indices, the establishment panel, and the county-year dispensing margin.

A.1 Category Definitions (Numerator)

Numerator departments are consolidated into estimation categories. Non-item departments (services, fees, sales tax, coupons, lottery, prepared-food/LSR–FSR nodes, and internal/testing nodes) are classified `non_item` and dropped. The household essentials category encompasses the non-grocery, non-prescription retail assortment; together with fresh and packaged groceries, it constitutes what the paper calls CPG.

1. **Other Household Essentials**: ATV & Powersport Accessories, ATV, Motorcycle & RV Tires, Abrasive & Finishing Products, Academic Journals & Periodicals, Accessories & Supplies (Electronics), Accessories (Sports), Action & Extreme Sports Gear, Action Figures & Figurines, Adult Literature & Fiction, Adult Romance Books, Air Fresheners & Deodorizers, Apparel Accessories, Architecture & Design Books, Arts & Crafts, Arts & Photography Books, Auto Safety & Security, Automotive Electronics, Automotive Interior Accessories, Automotive Replacement Parts, Baby Apparel, Baby Health Supplies, Balanced Nutrition Drinks & Snacks, Balloons, Banners, Streamers & Confetti, Bath & Body Care, Bathing & Skin Care (Baby), Bathroom Products, Batteries, Bedding & Decor (Baby), Bedding (Home & Garden), Binders & Binding Systems, Biography & Memoirs, Book Covers, Boys Apparel, Buckets, Building (Tools & Home Improvement), Building (Toys), Business & Investment Books, Business Office Furniture, Cake Supplies, Calendars, Camcorders & Accessories, Cameras & Camera Supplies, Cannabis Beverages, Car Care & Maintenance, Car Electronics, Cargo Storage & Racks, Carrying Cases, Caulk & Caulk Guns, Cell Phones, Childrens & Teen Books, Christmas Decorations, Cleaning Products (Baby), Cleaning Tools, Clock Radios & Alarm Clocks, Combustible Nicotine Products, Computers & Accessories, Computing & Internet (Books), Dance & Gymnastics Gear, Decor, Deodorants & Antiperspirants, Desk Accessories & Workspace Organizers, Diapering, Dishwashing, Dolls & Accessories, Drains & Gutters, Drones & Accessories, Ear, Easter Decorations, Education & Crafts, Educational Toys, Electrical, Electronics (Office), Electronics (Toys), Equipment (Baby), Everyday Decorations, Exercise & Fitness Gear, Exterior Automotive Accessories, Eye & Vision Care, FSR Beverages, Feminine Care, Filing Products, Flooring, Floral, Food & Beverage Cookbooks, Foot Hygiene & Foot Care, Footwear, Fragrance, Funeral & Memorial Products, Furniture (Home & Garden), GPS & Navigation, Games & Puzzles, Gardening & Lawn Care, Gasoline Cans, Gift Bags &

Wrapping Paper, Gift Baskets, Gift Cards, Gift Sets (Health & Beauty), Gifts (Baby), Girls Apparel, Hair Care, Halloween Decorations, Hand Hygiene, Hand Trucks, Hardware, Health & Beauty Tools, Health, Mind & Body Books, Heating, Cooling & Air Quality, Heavy Equipment & Agricultural Supplies, History Books, Hobbies, Home - Indoor / Game Room, Games & Accessories, Home - Outdoor Sports, Games & Accessories, Home Appliance Parts & Accessories, Home Audio & Theater, Home Telephones and Two-Way Radios, Home, Hobbies, & Garden Books, Household Cleaners, Infant & Toddler Nutrition, Installation Services, Invitations & Cards, Janitorial (Household), Janitorial Supplies (Office), Kids & Teens Rooms, Kitchen & Dining, Kitchen & Laundry Fixtures, Laundry, Leisure Sports, Games & Accessories, Lighters, Lighting & Ceiling Fans, Luggage & Travel Accessories, Magazines, Major Appliances, Makeup, Massage, Therapies & Relaxation, Medical Products, Mens Apparel, Mobile Accessories, Mobile App Downloads, Movies & Television Shows, Music (Entertainment), Music (Toys), Music Books, Musical Instruments, Newspapers, Noisemakers, Novelty (Toys), Novelty Apparel & Costumes, Nursery Furniture, Occupational Health & Safety Equipment, Office Equipment, Office Furniture, Office Furniture & Lighting, Office Lighting, Office Storage Supplies, Office Supplies, Office Supplies & Technology, Oral Hygiene, Other Office Equipment & Supplies, Outdoor Equipment & Accessories, Outdoor Power Equipment, PRRPs/NGPs, Painting Supplies & Wall Treatments, Paper & Plastic (Household), Parenting & Family Dynamic Books, Party & Occasions Table Accessories, Party Supplies, Patio Decorations & Furniture, Performance Nutrition, Performing Arts Books, Personal Health Care, Personal and Home Audio Devices, Pest & Insect Control (Household), Pet Food & Treats, Pet Prescriptions, Pet Supplies, Photo Products, Photography Books, Political & Social Science Books, Potty, Pregnancy & Postpartum Care Devices, Preschool Toys, Pressure & Temperature Gauges and Regulators, Racquet Sports Equipment, Reference Books, Religious Books & Texts, Restaurant Supplies, Rough Plumbing, Safety & Security, Safety (Baby), Science & Technology Books, Security & Surveillance Equipment, Sewing & Mending Supplies, Sexual Wellness Products, Shaving, Grooming, & Hair Removal, Shop Towels, Skin Care, Small Appliances, Smart Home, Smokeless Nicotine Products, Sports & Outdoor Play (Toys), Sports & Recreation Books, Stationery (Baby), Storage & Organization, Store Signs & Displays, Stuffed & Plush, Surface Care & Protection, Sustainability Books, TV & Video, Tablets & eReaders, Team Sports Equipment, Textbooks, Thanksgiving Decorations, Threading Taps & Dies, Tickets, Tires, Toddler Apparel, Toddler Furniture, Tools (Automotive), Tools (Tools & Home Improvement), Top Gifts, Towing & Winching, Toy Vehicles, Toys (Baby), Travel & Nature (Books), Ungendered Apparel, Utility Vehicles Parts & Accessories, Vacuums, Floor Care, & Accessories, Video Game Strategy Guides, Video Games, Consoles, & Accessories, Virtual Reality, Vitamins & Supplements, Water Quality Testing, Wearable Technology, Wedding Decorations, Weight

Management (Diet), Welding Equipment and Supplies, Womens Apparel, Work Gloves, Work Wear & Uniforms, iPods & MP3 Players.

2. **Packaged Grocery**: Alcohol Beverages, Baking & Cooking (Grocery), Beverages, Breakfast Foods, Candy (Snacks), Canned Foods, Condiments, Dairy, Dry Beans & Grains, Frozen Foods, Herbs & Spices, Ice, Packaged Bakery (Bread & Alternatives), Pasta & Noodles, Shelf Stable Meals, Snacks.

3. **Fresh Grocery**: Deli & Prepared Foods, Fresh Meal Kits, Fresh Meat, Fresh Produce, Fresh Seafood & Fish, In-Store Bakery (Bread & Alternatives), Misc. Food, Refrigerated Foods.

4. **Prescription Drugs**: Prescription (RX).

A.2 Store-Category Price Indices

The demand system is estimated on store-category price indices p_j^k rather than on individual products. Appendix B.1 offers the theoretical treatment: under weak separability and homotheticity, each continuous category admits a price index ρ_k, homogeneous of degree one in product prices, such that the category quantity index is deflated expenditure, $q^k = e_k/\rho_k$. This section describes how the empirical counterparts are built from the household panel. Store-category price indices of this kind also underlie the estimation in Thomassen et al. (2017); Atkin et al. (2018) discuss price-index construction from household purchase data.

Unit prices and aggregation. The inputs are the panel's item-level purchase records: the amount paid, the per-item price, and, where recorded, the package size. Prices are standardized by package size, so that unit prices are per equivalent unit and different pack sizes of the same product are comparable. Prescription items are excluded, and the remaining items are assigned to fresh groceries, packaged groceries, and other household essentials using the category mapping of Appendix A.1. Observations are then aggregated to UPC-by-banner-by-week cells, each carrying the median unit price and projection-weighted expenditure. The store dimension is the banner for named chains and a pooled composite banner for unbranded formats (independent pharmacies, dollar stores, other food stores, other retail)—the granularity at which stores enter the demand system; because the major chains priced nationally over the sample, the index carries no geographic dimension within a banner.

Hedonic regression. For each category k, a weighted least-squares regression with expenditure weights,

$$\log p_{ujt} \;=\; \gamma_u \;+\; \delta_{jt}^k \;+\; e_{ujt}, \tag{34}$$

projects log unit prices on product (UPC) fixed effects γ_u and banner-week fixed effects δ_{jt}^k, where u indexes products, j banners, and t year-weeks. The product fixed effects absorb product identity and quality, so δ_{jt}^k is identified from like-for-like comparisons of the same products sold across different banners and weeks. Identification requires overlap in assortments across banners, which is plentiful for packaged products. The category price index is the exponentiated fixed effect normalized to a base banner-week, $p_{jt}^k = \exp\left(\delta_{jt}^k - \delta_0^k\right)$. Expenditure weighting keeps the index representative of the category basket as actually purchased, and estimating (34) separately by category allows a banner to be relatively cheap in groceries yet expensive in household essentials, or vice versa.

Missing cells, availability, and quantities. Banner-category-weeks with no transactions are filled by log-linear interpolation within banner-category, with flat extrapolation at the sample ends. Banner-categories that never transact are treated as unavailable rather than missing—fresh groceries at chain pharmacies is the leading case—which is the availability structure the demand model exploits (Section 4). The same index defines quantities in estimation: observed category expenditure at a store is deflated by the store's index, $q^k = e^k/p_j^k$, so price and quantity units are internally consistent, and the base normalization is simply a choice of quantity units. The drug price index is set to one at all pharmacies: the prescription price is payer-set, exhibits no usable cross-sectional variation, and enters demand without a price term (Section 4); Appendix B.3 separates the exogeneity and retailer-invariance content of this treatment.

Price instrument. Alongside each index, I construct the leave-one-out average of the other banners' indices in the same category and week for every banner-category-week. This is the Hausman-style instrument (Hausman et al., 1994) constructed as a safeguard for the price moments; the baseline estimates rest on the conditional-independence assumption of Section 5, and re-estimation with the instrument is a planned robustness exercise. Rival indices co-move with a banner's own index through common cost shocks—wholesale acquisition and distribution costs—while excluding banner-specific demand shocks; the residual threat — demand shocks common to all banners within a week —is absorbed by the time controls in Section 5.

A.3 The Establishment Panel

The establishment panel treats the physical site as the economic unit. It is built from the business registry's clustered establishment records, whose cross-year identifiers are stable to name and ownership changes. Operator types are assigned in two stages—named banners are mapped directly to types, and unbranded records are classified by industry code (Table A1)—after which each company is assigned its modal type across sites and years, eliminating spurious type churn from year-to-year industry-code revisions.

Table A1: Store-Format Classification by NAICS Code

	Classification vintage		
Store concept	2012	2017	2022
Grocery and convenience			
Supermarkets and other grocery	44511	44511	44511
Convenience and vending	44512	44512	44513
Specialty food			
Meat retailers	44521	44521	44524
Fish and seafood	44522	44522	44525
Fruit and vegetable	44523	44523	44523
Other specialty food	44529	44529	44529
Beer, wine, and liquor			
Beer, wine, and liquor	44531	44531	44532
General merchandise			
Department stores	45211	45221	45511
Warehouse clubs and supercenters	45291	45231	45521
Health and personal care			
Pharmacies and drug stores	44611	44611	45611
Gasoline			
Gas stations with convenience	44711	44711	45711

Notes: Classification of unbranded establishment records into operator types by NAICS industry code and classification vintage. Named banners are mapped directly to types before the industry-code stage, and each company is then assigned its modal type across sites and years.

Activity spells are constructed to measure economic entry and exit rather than data coverage. Interior gaps within an establishment's record are treated as coverage holes rather than exits, so a spell ends only at the final observed year. Ownership transfers are not exits: when an exiting establishment is succeeded by a same-type establishment at approximately the same location within a year, both the exit and the corresponding entry are voided and flagged as a transfer. The adjustment is quantitatively important—transfers account for roughly one quarter of independent exit records and roughly two-fifths of chain exit records, the latter concentrated in the 2018 divestiture of Rite Aid stores—and without it, the exit hazards feeding the choice probabilities would conflate closures with ownership changes. Establishment identifiers appearing only within a single registry vintage are dropped as ingestion artifacts, and left- or right-censored spells contribute no entry or exit events.

A.4 The County-Year Dispensing Margin

The county-year dispensing margin $m_{r,mt}$ combines public CMS data with the national NCPA benchmark of Section 2. Medicare monthly enrollment measures each county's payer mix (Part D, Medicare Advantage, dual-eligible, and low-income-subsidy shares); state Medicaid fee schedules report professional dispensing fees by state and quarter; and published aggregates of Part D pharmacy price concessions—direct and indirect remuneration (DIR) fees—capture the clawbacks that eroded Part D dispensing margins over the sample. Weighting the payer-specific margins by county payer mix, and anchoring the enrollment-weighted national mean to the NCPA per-prescription margin series $\bar{m}_{rt}$, yields the county-year margin; county-years without coverage carry the national value and contribute no cross-sectional variation. The margin enters the supply model as a state variable, and its decomposition into the national path and the county-level payer deviation is the source of identifying variation for the dollar scale (Section 7.5).

B Demand Model Appendix

This appendix presents the details of the demand model: the microfoundations of the estimation specification, its relation to Thomassen et al. (2017), the evidence for the constant-price treatment of prescriptions, and the estimation details.

B.1 Microfoundations and First-Order Conditions

This section derives the estimating model of Section 4 from a general store-and-quantity choice problem, following Appendices D and E of Thomassen et al. (2017), and states the first-order conditions on which the estimation rests. Consumer and time subscripts are suppressed.

B.1.1 The General Problem

Let $\mathcal{J}$ denote the consumer's store choice set. In the most general formulation, the consumer chooses non-negative quantities of every product at every store, a fill decision, and an outside good q_0 (price normalized to one), subject to a budget constraint:

$$\max_{\mathbf{q}\geq 0,\ y\in\{0,1\},\ q_0\geq 0} V(\mathbf{q}, y, q_0) \qquad \text{s.t.} \qquad \mathbf{p}'\mathbf{q} + p_r\, y + q_0 \leq \text{inc},$$

where $\mathbf{q}$ stacks product-level quantities over all stores in $\mathcal{J}$ and inc is income. Corner solutions arise from two distinct sources—stores the consumer does not visit, and products she does not buy at visited stores—and a demand system over all store-product quantities with corners of both kinds is not estimable in practice. The estimating model results from the following assumptions.

B.1.2 Assumptions

Each of the four assumptions plays a distinct role:

A1. (*Shopping structure*) The consumer visits at most two stores per occasion and purchases each category at a single store. Both restrictions are read off the data (roughly 92 and 96 percent of occasions, respectively; Table 5) rather than imposed *a priori*.[35] The problem then separates into a discrete basket choice c, a store-per-category assignment d, and quantity choices, and can be written

$$\max_{c\in\mathcal{C}}\ \max_{d}\ \max_{\mathbf{q}\geq 0,\, y}\ V(c, d, \mathbf{q}, y, q_0) \qquad \text{s.t.} \qquad \mathbf{p}_d'\mathbf{q} + p_r\, y + q_0 \leq \text{inc.}$$

A2. (*Separable shopping costs*) Utility is additively separable between the variable utility of the basket and the shopping cost: $U = V(c, d, \mathbf{q}, y, q_0) - \Gamma(c) + \varepsilon_c$. Shopping costs then shift *which* baskets are chosen but not the marginal rates of substitution among goods within a basket—the property that allows the discrete and continuous layers to be estimated jointly yet identified from distinct variation.

A3. (*Quasilinearity*) V is linear in the outside good, with coefficient α. Substituting the budget constraint eliminates income from the within-basket problem: category demands have no income effects, and income influences behavior only through the price-sensitivity coefficient α_i, where it is modeled directly. For the small share of weekly expenditure at stake on a shopping occasion, this is a standard and defensible abstraction.

A4. (*Weak separability and homotheticity*) Preferences over products are weakly separable across the categories $\{x, f, p, y\}$, and the subutility of each continuous category is homothetic. Weak separability delivers two-stage budgeting, and homotheticity makes the within-category stage summarizable by scalars, as derived next.

B.1.3 Two-Stage Budgeting and Category Indices

Under A4, utility can be written over category subutilities, $u = U\big[v_x(\mathbf{x}_x),\ v_f(\mathbf{x}_f),\ v_p(\mathbf{x}_p),\ y\big]$, and the consumer's problem decomposes into an upper stage that allocates expenditure across categories and lower stages that allocate category expenditure across products. For a continuous category k with product prices $\mathbf{p}_k$, the lower-stage indirect utility is $\psi_k(e_k, \mathbf{p}_k) = \max\{v_k(\mathbf{x}_k) : \mathbf{p}_k'\mathbf{x}_k = e_k\}$; homotheticity implies

$$\psi_k(e_k, \mathbf{p}_k) = e_k/\rho_k(\mathbf{p}_k),$$

with ρ_k homogeneous of degree one. The scalar $q^k \equiv e_k/\rho_k(\mathbf{p}_k)$ serves as the category quantity index and ρ_k as the category price index, so the upper stage is a choice over (q^x, q^f, q^p, y, q_0) facing prices $(\rho_x, \rho_f, \rho_p, p_r, 1)$.

[35] The single-store-per-category restriction could in principle be relaxed by doubling the quantity vector when $n(c) = 2$ and adding second-order parameters that govern within-category, cross-store substitution; given how little consumers spend on the same category at a second store, the restriction is innocuous here.

This is the theoretical basis for estimating demand using store-category price indices rather than individual products; the empirical indices are constructed at the banner-category level, as described in Appendix A.2. The drug requires no lower stage: it is a unit-demand category whose "index" is the fill indicator itself. Under additive (strong) separability across categories, homotheticity can be replaced by the weaker requirement that product-level subutilities take a Generalized Gorman Polar Form (Thomassen et al., 2017, Appendix E).

B.1.4 Quadratic Specialization

Specializing the upper-stage utility to a second-order form in (q, y), where $q = (q^x, q^f, q^p)'$, and imposing A3:

$$U = \mu' q + \beta_r\, y - \tfrac{1}{2} \begin{pmatrix} q \\ y \end{pmatrix}' \begin{pmatrix} \Lambda & \lambda_{\cdot y} \\ \lambda_{\cdot y}' & \lambda_{yy} \end{pmatrix} \begin{pmatrix} q \\ y \end{pmatrix} + \alpha q_0, \qquad \lambda_{\cdot y} = (\lambda_{xy}, \lambda_{fy}, \lambda_{py})'.$$

Substituting the budget constraint for q_0, dropping the constant $\alpha\,\mathrm{inc}$, using $y^2 = y$ to absorb $-\frac{1}{2}\lambda_{yy}\, y$ and $-\alpha p_r\, y$ into the linear drug value, and appending the shopping-cost and basket-shock terms from A2 yields (4).

With the store dimension discretized and each category reduced to a quantity index, the quadratic is a second-order flexible form for the upper-level utility over category indices: the linear terms carry store-category quality and prices, the diagonal curvature generates satiation, and the off-diagonals are exactly the cross-category interactions the paper is about. Two properties recommend the quadratic over familiar alternatives. First, it natively accommodates zero demands: budget-share systems such as the Almost Ideal Demand System (Deaton and Muellbauer, 1980) are ill-suited to corners, whereas quadratic utility with nonnegativity constraints has served this purpose since Wales and Woodland (1983). Second, its first-order conditions are linear on any candidate set of purchased goods, which the estimator of Section 5 exploits heavily. The binary drug fits the structure as a unit-demand category: its "quantity index" is the fill indicator, its within-category stage is trivial, and its own curvature is absorbed into the linear drug value.

B.2 Sources of the Drug–Retail Linkage and Relation to Thomassen et al. (2017)

It is worth being precise about how the model generates the drug–retail linkage and what is inherited from Thomassen et al. (2017). The model contains two conceptually distinct sources of cross-category linkage. The first is structural complementarity—the off-diagonal λ's—by which a fill mechanically shifts the marginal utility of retail goods on the same occasion, raising it when the interaction is negative. The second is the shopping-cost channel: because visiting a second store is costly, consumers concentrate purchases, so anything that brings a consumer into a store (a prescription need) exposes other categories to that store's prices and quality. Either channel can make pharmacies attract retail spending. However, they differ in policy content: the first is a preference primitive, while the second responds to the choice set and prices. The model identifies them separately: the λ's from fill-conditional quantity gaps within stores and the shopping-cost parameters from discrete basket choice. Correlated tastes—a consumer who likes shopping buys more of everything on the same trip—mimic complementarity in a cross-section, and the panel moments of Section 5 exist to absorb this correlation into σ_1, σ_2 before the λ's are read off cross-category comovement, following the logic of Gentzkow (2007).

Relative to Thomassen et al. (2017), the utility specification (4)–(5), the two-store basket structure, the corner-solution treatment, the heterogeneity specification (12)–(14), and the moment-based estimation strategy are direct descendants of their equations for supermarket demand; the crosswalk is deliberately tight so that the departures are interpretable. The departures are the binary, price-regulated drug category and its deterministic-fill regime; availability gating of fresh groceries across formats; the pharmacy-format intercepts identified from fill location; and the object of interest itself—not the level of cross-category effects and their implication for static market power, but the *time path* of the drug–retail complementarity and its implication for which store formats can survive.

B.3 Out-of-Pocket Prices and the Constant-Price Treatment of Prescriptions

The demand system carries no prescription price term: the drug price index is one at every pharmacy (Appendix A.2), fills are allocated across formats by the drug intercepts $\xi_{b,y}$ (Section 4.4), and the consumer cost of a closure is travel and complementarity, not a price change. This treatment bundles two logically distinct claims. The first is *exogeneity*: the out-of-pocket price is set by plan design—copay tiers, coinsurance rates, and deductibles emerging from contracts among payers, PBMs, and employers—and is not a choice variable of the pharmacy, so there is no drug analog of the retail price endogeneity problem of Section 5. The second is *retailer-invariance*: a given consumer filling a given prescription faces the same out-of-pocket price at every pharmacy in the choice set. Identification uses the first claim, and the constant-price treatment uses the second.

Medicare Part D preferred networks. An objection may come from Medicare Part D, where preferred cost-sharing is now essentially universal among standalone prescription drug plans, and preferred status varies across chains—across-retailer variation in consumer prices, concentrated in the heavy-fill population most exposed to closures (Fein, 2024; Cubanski and Damico, 2023). The objection has less force than it first appears, for the absorption reason above: preferred-network differentials are set at the chain level and are stable within the plan year, so they are format-constant differentials of exactly the kind the drug intercepts absorb, shifting where fills land—which the intercepts are fitted to match—without contaminating the complementarity or travel parameters (Starc and Swanson, 2021). The channel they could affect is the welfare accounting, if closure-induced relocation systematically moved Part D fills from preferred to non-preferred pharmacies; that channel is bounded by the same arithmetic as the cash segment below, with the tier differential in place of the cash differential.

The cash segment and a bound on the welfare error. The genuine across-retailer dispersion in consumer drug prices lies in the cash-and-discount-card segment—direct-to-consumer cards such as GoodRx, retailer generic lists, and copayment coupons—where the price of the same generic can differ several-fold across retailers. The segment is small, on the order of five percent of dispensed prescriptions, and concentrated in low-cost generics (IQVIA, 2024)—unsurprisingly, since it exists where a cash price can undercut a copay. For the closure welfare statements, the omitted term is bounded by (cash-segment share of the closing store's fills) $\times$ (per-fill price differential between the closed store and the relocation target): even at a generous \$10 differential, the bound is a fraction of a dollar per relocated fill, second order against the travel and complementarity components of the removal CVs (33).

B.4 Estimation Details

Sample construction details. Beyond the screening steps of Table 4, occasions with implausibly long chosen shopping tours (over 60 miles) are dropped, and dominated alternatives beyond a distance cap (over 90 miles) are trimmed from the choice set. To make estimates comparable across years, quantity winsorization caps and demographic centering constants are computed once from the pooled 2019–2024 sample and frozen, so that parameters retain the same units in every year (Section 5.1).

Reparameterization and optimization. The optimizer searches over a reparameterized space that enforces the model's restrictions by construction: Λ is parameterized through its Cholesky factor, guaranteeing positive definiteness, with the drug cross-terms $\lambda_{xy}, \lambda_{fy}, \lambda_{py}$ left unconstrained since their sign is the object of interest; positivity of the scales, spreads, and price coefficients is imposed through exponential transforms; and the normalizations of Section 4 ($\beta_{0x} = 1$, $\beta_{0,y} = 0$, the reference-format and structural-zero drug and fresh intercepts) are held. Estimation uses a deterministic multi-start (24 starting points) with a quasi-Newton optimizer, and the delta method maps the reparameterized estimates and their covariance back to the structural parameters.

Simulation discipline. Simulation follows a strict discipline required by the panel moments: draws are fixed across the optimization; a consumer's adjacent occasions share the permanent draws; and the store-category shocks ν_{ijk} are drawn once per (consumer, store, category) and reused across all of the consumer's occasions, without which σ_4 would not be identified. The corner-aware quantity problem is solved by an exact active-set enumeration implemented in a compiled kernel and verified against an independent reference implementation.

C Supply and Counterfactual Appendix

This appendix presents details of the supply side: the uniform-reimbursement assumption; the derivation of the estimating equations from the conditional choice probability inversion; the identification arguments and the estimation procedure for the dynamic model; and the design assumptions underlying the counterfactuals.

C.1 (Near-)Uniform Reimbursement

The treatment of reimbursement as (almost) uniform across the pharmacies of a county-year (Section 7.2) is a substantive assumption about a market in which reimbursement is negotiated between pharmacies and PBMs. Its principal threat is vertical integration: the largest PBMs are integrated with chain pharmacies—CVS–Caremark being the leading case—and an integrated PBM could in principle reimburse rival pharmacies, independents especially, less generously than its own stores. This is exactly the margin on which Yde (2025) sheds light.

Three considerations govern how much the caveat matters here. First, the direct evidence is contested: an independent audit commissioned in response to these concerns finds that CVS Caremark does not reimburse independent pharmacies at rates lower than those for its affiliated stores (CVS Health, 2024). Second, the regulatory environment has shifted: the Federal Trade Commission maintains an active inquiry into PBM practices and has signaled renewed interest in Robinson–Patman enforcement against discriminatory pricing (Federal Trade Commission, 2024), so a study period such as Yde's may well capture an era of discriminatory reimbursement that present-day scrutiny makes far less likely. Third, and most importantly, uniformity does not require that the *level* be competitive: nothing here rules out that vertically integrated payers set the uniform rate in their own favor—pricing power exercised subject to uniform application—and that level is precisely the object the counterfactuals move. The question of how integration affects drug prices is deliberately left to the studies built to measure it; this paper's mechanism is that closures are driven by the erosion of front-store retail—which is exactly why the integrated chains are the ones closing stores—and that mechanism is orthogonal to the integration debate.

C.2 The Hotz–Miller Inversion: Deriving the Log-Odds Representations

I derive the estimating equations (26) and (27) of Section 7.4 here. The incumbent equation follows from the Hotz–Miller inversion (Hotz and Miller, 1993) together with the terminal-action property of exit; the entry equation requires one additional step—the renewal argument of Arcidiacono and Miller (2011)—because staying out of a slot, unlike exiting a store, is not a terminal action.

Index units. Fix a site s. Conditional expectations are with respect to $x_{s,t+1}$ given x_{st}, and the size arguments k_j, k_s are suppressed throughout without loss of generality. Payoff shocks are i.i.d. type-I extreme value with dollar scale σ_ε^τ (Section 7.3). Dividing every payoff by σ_ε^τ converts the problem to *index units*—one dollar becomes $\psi^\tau = 1/\sigma_\varepsilon^\tau$ index units, per (25)—and leaves unit-scale shocks. All value objects below are in index units. This is why the continuation corrections in (26)–(27), being built from log-probabilities, enter with coefficient exactly β and no ψ^τ attached: they are constructed in shock-scale units.[36]

Two properties of extreme-value choice. Consider any discrete choice over alternatives $b \in \mathcal{B}$ with systematic values $v_b(x)$ in index units and additive i.i.d. unit-scale type-I extreme-value shocks, so that choice probabilities take the multinomial logit form $P(b \mid x) = e^{v_b(x)}/\sum_{b'} e^{v_{b'}(x)}$. Two exact identities follow. First, the *inversion*: for any pair of alternatives,

$$v_b(x) - v_{b'}(x) \;=\; \ln P(b \mid x) - \ln P(b' \mid x), \tag{35}$$

so differences in conditional values are point-identified by observed choice probabilities (Hotz and Miller, 1993). Second, the ex-ante value $V(x) \equiv \mathbb{E}\max_b\{v_b(x) + \varepsilon_b\}$ satisfies, for *every* alternative b,

$$V(x) \;=\; \ln \sum\nolimits_{b'} e^{v_{b'}(x)} + \gamma_{\mathrm{E}} \;=\; v_b(x) + \gamma_{\mathrm{E}} - \ln P(b \mid x), \tag{36}$$

where $\gamma_{\mathrm{E}} \approx 0.577$ is the Euler–Mascheroni constant (unrelated to the site-cost loadings γ). Identity (36) is the engine of the collapse: the ex-ante value can be written relative to *any single* alternative, with $\gamma_{\mathrm{E}} - \ln P(b \mid x)$ correcting for the option value of the alternatives not taken. If some alternative has a *known* conditional value—because the action terminates the problem or returns the state to a common position—then (36) collapses the entire continuation value to a one-period-ahead log-probability, with no further recursion.

[36] The entry nest mixes formats with different dollar scales. The primitive is stated directly in index units: each alternative's systematic index converts that format's dollars at the format's own rate ψ^a—so a dollar of type-a profit moves the entry index of a exactly as it moves the stay index of a type-a incumbent, which is what the counterfactual uses when a margin change shifts the entry index by $\psi^a\,\Delta m\,Y^a$—and the shocks are i.i.d. unit scale across alternatives, so the multinomial logit form and the inversion below apply within the nest. The staying-out alternative carries no flow dollars, so no conversion rate attaches to it.

The incumbent equation. An active type-τ occupant chooses between staying and exiting. In index units, the conditional values are

$$v_S^\tau(x_{st}) \;=\; \psi^\tau\big[\pi^{*,\tau}(x_{st}) - \phi^\tau(s)\big] \;+\; \beta\,\mathbb{E}\big[V^\tau(x_{s,t+1})\,\big|\,x_{st}\big], \qquad v_X^\tau(x_{st}) \;=\; 0, \tag{37}$$

where V^τ is the incumbent's ex-ante value. The exit value is zero because exit is terminal for the establishment, and the scrap value is normalized to zero: the vacated slot returns to the entrant pool, but its subsequent value accrues to future entrants rather than to the exiting firm. Exit is therefore the known-value alternative. Applying (36) at $t+1$ with $b = X$,

$$V^\tau(x_{s,t+1}) \;=\; v_X^\tau(x_{s,t+1}) + \gamma_{\mathrm{E}} - \ln P_\tau^X(x_{s,t+1}) \;=\; \gamma_{\mathrm{E}} - \ln P_\tau^X(x_{s,t+1}), \tag{38}$$

and substituting into (37) and applying the inversion (35) to the stay/exit pair gives the log-odds

$$\ell_\tau^X(x_{st}) \;=\; v_S^\tau(x_{st}) - v_X^\tau(x_{st}) \;=\; \psi^\tau\big[\pi^{*,\tau}(x_{st}) - \phi^\tau(s)\big] \;+\; \beta\gamma_{\mathrm{E}} \;+\; \beta\,\mathbb{E}\big[-\ln P_\tau^X(x_{s,t+1})\,\big|\,x_{st}\big]. \tag{39}$$

The constant $\beta\gamma_{\mathrm{E}}$ is absorbed into the freely varying fixed-cost intercept $\psi^\tau\phi_0^\tau(m)$ (equivalently, into the fixed effects of the estimating regressions), which yields (26). The correction $-\ln P_\tau^X(x_{s,t+1}) \geq 0$ is the value of arriving at $t+1$ as an incumbent up to the constant—large where exit is unlikely, that is, where continuing is valuable.

The entry equation. A potential entrant at an open slot chooses $a \in \{\varnothing, C, I, P\}$. In index units, entering as a strategic format yields

$$u_a(x_{st}) \;=\; \psi^a\big[\pi^{*,a}(x_{st}) - \phi^a(s) - \kappa^{E,a}\big] \;+\; \beta\,\mathbb{E}\big[V^a(x_{s,t+1})\,\big|\,x_{st}\big], \qquad a \in \{C, I\}, \tag{40}$$

while staying out yields no flow and leaves the option standing at the same slot,

$$u_\varnothing(x_{st}) \;=\; \beta\,\mathbb{E}\big[U(x_{s,t+1})\,\big|\,x_{st}\big], \tag{41}$$

where U is the ex-ante value of the potential entrant. Staying out is *not* a terminal action—$u_\varnothing$ is an unknown function of the state—so the collapse (38) does not apply directly. What replaces terminality is finite dependence: *exit is a renewal action for the slot.* Compare two action paths from the same state x_{st}: (i) enter as a at t, exit at $t+1$; (ii) stay out at t, stay out again at $t+1$. Under path (i), the slot is vacated during $t+1$ and re-enters the entrant pool at $t+2$; under path (ii), it remains vacant and is again at risk at

$t+2$. Both paths thus deposit the slot in the same position two periods out—vacant and at risk—facing the same conditional distribution of $x_{s,t+2}$, because a single site's own action does not move the law of motion of its features: the distance bands exclude the site itself, and rivals' responses are expectational innovations under rational expectations (Section 7.4). The value of that common terminal position cancels from the comparison, and only flow payoffs and one-period-ahead corrections survive.

Formally, apply (36) at $t+1$ along each path, choosing as reference the action the path prescribes. Along the entry path, with $b = X$,

$$\mathbb{E}\big[V^a(x_{s,t+1}) \,\big|\, x_{st}\big] \;=\; \mathbb{E}\big[v^a_X(x_{s,t+1}) + \gamma_{\mathrm{E}} - \ln P^X_a(x_{s,t+1}) \,\big|\, x_{st}\big]; \tag{42}$$

along the waiting path, with $b = \varnothing$,

$$\mathbb{E}\big[U(x_{s,t+1}) \,\big|\, x_{st}\big] \;=\; \mathbb{E}\big[u_\varnothing(x_{s,t+1}) + \gamma_{\mathrm{E}} - \ln P^E_\varnothing(x_{s,t+1}) \,\big|\, x_{st}\big]. \tag{43}$$

The leading terms are the values of the positions in which the two paths leave the slot at the end of $t+1$: an establishment that has just exited and a claimant that has just waited hold the same object, a vacant slot that will be at risk at $t+2$, so

$$v^a_X(x_{s,t+1}) \;=\; u_\varnothing(x_{s,t+1}) \tag{44}$$

is the renewal condition: exit undoes entry. Subtracting (41) from (40) using (42)–(44), the position values and the constants γ_{E} cancel,

$$u_a(x_{st}) - u_\varnothing(x_{st}) \;=\; \psi^a\big[\pi^{*,a}(x_{st}) - \phi^a(s) - \kappa^{E,a}\big] \;+\; \beta\,\mathbb{E}\big[\ln P^E_\varnothing(x_{s,t+1}) - \ln P^X_a(x_{s,t+1}) \,\big|\, x_{st}\big], \tag{45}$$

and the inversion (35) applied within the entry nest equates the left side to $\ln\big[P^E_a(x_{st})/P^E_\varnothing(x_{st})\big]$, which is (27). No constant needs to be absorbed here, unlike in the incumbent equation. The interpretation matches the main text: $-\ln P^X_a(x_{s,t+1})$ is, up to the constant, the post-entry continuation value, and $-\ln P^E_\varnothing(x_{s,t+1})$ is the option value of the vacant slot, which entering today forfeits. The dollar-and-convenience alternative participates only through the denominator of the multinomial probabilities—it shifts the levels of P^E_a and $P^E_\varnothing$ but not the pairwise inversion—so no structural parameter attaches to it.

The derivation also makes precise the main text's remark that the binary occupancy structure licenses the collapse: were entry bundled with a persistent sub-choice (a quality investment, say), the enter-then-exit path would leave the slot in a different position than the wait-then-wait path, the renewal condition (44) would fail, and continuation values would survive the differencing. In estimation, the conditional expectations in (39) and (45) are replaced by fitted choice probabilities at the realized next-period state, as described in Section 7.4; nothing in the derivation changes.

C.3 Identification Details

The re-occupancy margin. The site formulation adds one identification source with no count-based analog. When a pharmacy exits and a dollar store occupies the same premises within a year or two, comparing the location's predicted CPG-only profitability with that of other dollar-store entry sites—within county, holding the site fixed—isolates the format contrast from location quality. This is the supply-side mirror of the paper's demand-side mechanism: if the CPG business at pharmacy corners remains viable. In contrast, the dispensing business does not; CPG-only formats should disproportionately claim vacated pharmacy sites, and their predicted CPG profitability should not lag that of other dollar-store sites. The re-occupancy share and the within-site profitability comparison are left as robustness checks to the structural estimates.

Universe censoring and the at-risk convention. Because the site universe is censored at ever-occupied locations, the entry risk set requires a convention: a site contributes vacancy years from the start of the panel (so that first entries are not truncated out of the sample) or only after its first observed occupancy (so that pre-history vacancy is not imputed). I use the panel-start convention. The convention matters for the level of the entry probability—and hence interacts with the level of κ_u^τ.

C.4 Estimation Procedure and Inference

Estimation proceeds in the order dictated by the layering of Section 7.5.

(i) Site panel. Sites, capacities, occupancy spells, and at-risk slots are constructed from the establishment panel of Appendix A.3; by construction, site occupancy aggregated to county-year-type counts reproduces the observed market histories exactly.

(ii) First stages. Exit probabilities are type-specific logits and the entry probability is a multinomial logit over $\{\varnothing, C, I, P\}$, both on the local features x_{st}. Their conditioning set includes the reimbursement margin $m_{r,mt}$ and year effects, and this is not optional: the scale regression works with fitted log-odds, so if the CCPs were not permitted to vary with the margin, the payer-driven component of profit would be orthogonal to the left-hand side by construction and ψ^τ would be attenuated to zero mechanically. The large-box states enter as observed covariates rather than through an estimated transition process, and the store-level profit projections are fit to the matched establishment profits and evaluated at every site-year and type (occupied or vacant).

(iii) Entry costs. Entry costs κ_u^τ are estimated from the within-cell contrast for $\tau \in \{C, I\}$.

(iv) Scale, fixed costs, and site-cost loadings. The offset regression (30) recovers them.

Inference is layered to match the identifying variation. For κ_u^τ, a county-block bootstrap resamples entire market histories and re-estimates every first stage—profit projections and choice probabilities—within each draw, so that reported standard errors carry demand-projection and CCP estimation error. For ψ^τ, whose identifying shocks are state-year fee changes, standard errors cluster at the state level, and the bootstrap resamples states as blocks; treating counties as independent draws would overstate precision, since a single fee change moves every county in a state. Due to computational capacity constraints, standard errors are not yet bootstrapped and only incorporate the inner-loop variation.

C.5 Counterfactual Design: Beliefs and Limitations

Beliefs and credibility. Entry and exit are forward-looking: what moves a store's decision is the expected discounted stream of margins, not the current margin alone. Within the model, an actual increase today and a credible commitment to an equal permanent increase are close substitutes. The margin enters both current flow profit and, through beliefs about its future path, the continuation value, so a fully believed announcement that margins will be permanently higher from some date forward moves entry and exit *before a dollar is spent*. In contrast, an increase paid today but believed to be temporary buys far less than the stationary comparison would imply. The counterfactual numbers should accordingly be read as the response to a fully believed shift of the margin *process*—however that belief is created—with the fiscal cost computed for the implementation in which the higher margin is actually paid on every fill from the first period onward.

This near-equivalence is exactly why credibility cannot be assumed for free. A payer that could move exit behavior with announcements alone has an incentive to renege once the entry response has materialized, and the DIR-fee experience gives pharmacies concrete reasons to discount reimbursement promises. To the extent operators discount permanence, the reported count responses are upper bounds on what a dollar of enacted policy achieves.

Long-run comparison. Comparing invariant distributions abstracts from the transition path: the count response accumulates over several years of entry and averted exit, while the outlay begins immediately. Short-run cost-effectiveness is therefore weaker than the stationary numbers suggest.

The environment is today's. Demographics, the large-box pool, and retail profitability are frozen at their current configuration (CPG-only formats churn at their fitted baseline process), so both long-run configurations are resting points at *current* fundamentals. Reimbursement can offset but not undo the retail-side mechanism that Section 3 identifies as the closure driver; counterfactual prices that offset in today's retail environment, together with continued CPG decline, would shift both configurations downward.

Drug prices are unchanged; retail prices re-optimize. A uniform margin change is a payment between payers and pharmacies: copays and drug prices are payer-set and held fixed. I do not model the premium or tax channel that finances the outlay, so the welfare statements are transfer accounting with a consumer-surplus term rather than a full incidence analysis. Retail shelf prices, by contrast, are re-solved under each configuration by the static re-pricing of Section 8.2, and the price channel enters the welfare calculation. The entry/exit simulation itself still evaluates flow profits at baseline prices; the fully iterated version—each period's configuration re-priced and fed back into the exit and entry indices—is deferred, and its omission slightly overstates the baseline decline, since re-pricing raises survivor margins.[37] On the drug side, I maintain the assumption that a consumer who relocates a fill after a closure faces no out-of-pocket consequence, so the consumer cost of a lost store is travel, availability, and retail prices. Appendix B.3 documents why this is the right first order and bounds the error from the cash-pay segment, where prices do vary across retailers.

[37] Still, the price level has always been higher for the chain pharmacies, and given an increased level of price sensitivity in real terms, I expect that the re-pricing feedback effect would only marginally salvage the baseline decline.