Health Insurance Design and the Passthrough of List Prices to Out-of-Pocket Costs*

Josh Feng, Luca Maini†

June 30, 2026

Abstract

The cost of prescription drugs to U.S. health plans depends on nominal "list" prices and confidential rebates. As rebates have grown in recent years, list prices have become increasingly detached from actual acquisition costs, prompting speculation about what is driving the divergence. Using a model of health insurance design, we argue that plans facing minimum coverage regulation can increase patient cost-sharing by covering drugs with higher list prices, thereby circumventing these restrictions. Using an instrumental variable approach, we confirm that higher list prices pass through to higher patient out-of-pocket costs, but only in plans whose benefit design is constrained. Passthrough is near one-to-one in Medicare Part D plans, which must meet minimum benefit-design regulation, and in commercial high-deductible health plans, which are constrained by the structural 100% cost-sharing cap below the deductible. Conversely, and consistent with the model, passthrough is lowest in more generous PPO and HMO commercial plans that can set cost-sharing freely. Relative to commercial plans, Medicare Part D plans are also much more likely to cover high-list-price biologics than highly similar, low-list-price biosimilar drugs.

*We thank Leila Agha, Joe Levy, David Ridley, Mark Shepard, Amanda Starc, and Ben Vatter for providing detailed comments and feedback, as well as conference participants at ASHEcon, Bates White, the Boston Conference on Markets and Competition, and the Southeastern Health Economic Study Group. We thank Richard Evans and Scott Hinds at SSR Health for providing data access and additional feedback. Karishma D'Souza, Hailey Brace, and Katherine Garrett provided excellent research assistance. This research project was supported by Arnold Ventures LLC.

†Josh Feng: University of Utah Eccles School of Business, josh.feng@eccles.utah.edu. Luca Maini: Harvard Medical School, maini@hcp.med.harvard.edu.

"Cost-sharing is often based on the list price of a drug, even when insurance companies and PBMs are paying a fraction of that price. Our current system [...] incentivizes high list prices and large rebates as a mechanism to keep insurance premiums low."

Kenneth Frazier (Chairman and former CEO, Merck & Co, Inc.)

1 Introduction

Prescription drug prices in the United States have risen sharply in recent years, drawing sustained scrutiny from policymakers concerned about their impact on healthcare spending and patient access. Drug manufacturers usually respond by arguing that only nominal, "list" prices have grown, whereas net-of-rebate prices—a better reflection of actual acquisition costs—have remained relatively stable (Aitken et al., 2016; Kakani et al., 2020). Still, this divergence matters, because high list prices can translate into higher out-of-pocket costs for patients under the deductible or who owe coinsurance. As evidence mounts that higher cost-sharing can lead patients to discontinue high-value drugs (Brot-Goldberg et al., 2017) and even raise mortality (Chandra et al., 2024; Roberts et al., 2025), understanding why health plans would choose to tie cost sharing to an ever-rising price that is detached from the true cost of care is a question of increasing importance.

In this article, we argue that the divergence of list and net prices and the tying of list prices to cost sharing are driven by payer incentives created by regulatory constraints. We begin by showing that under a canonical model of a competitive insurance market with adverse selection (Rothschild and Stiglitz, 1976), health plans facing minimum coverage requirements in the form of a cap on coinsurance rates optimally choose to cover drugs with higher list prices as a way to increase out-of-pocket costs and lower premiums. Then, we test the implications of the theory using an instrumental variable design to estimate the passthrough rate of list price to out-of-pocket spending across different U.S. health insurance market segments and plan types. Consistent with our theory, we find estimates close to 1 for plans offered in the Medicare Part D marketplace, where government regulation mandates a minimum standard benefit design consisting of a set of coinsurance rates. Conversely, plans offered on the

employer-sponsored insurance market—which is largely unregulated—display virtually no passthrough of list prices to out-of-pocket costs.[1]

The intuition for the theoretical result is simple: by tying cost sharing to list prices and covering drugs with higher list prices, plans can reduce their effective generosity while appearing more generous *on paper*—a strategy that is generally optimal in markets with adverse selection, or where patients care more about premium than expected out-of-pocket costs (e.g., as in Abaluck and Gruber, 2011). Crucially, however, this only happens in government-regulated markets where plans face coverage requirements on their nominal generosity, or in corner solutions where cost-sharing is mechanically capped (e.g., at 100% when under the deductible). Conversely, when plans face no constraints on their benefit design, list prices do not affect the outcome of the model because plans can increase out-of-pocket costs without the aid of list prices. A key implication of the model is therefore that list prices should only affect out-of-pocket costs for patients enrolled in such "constrained" plans.

As discussed in a previous paragraph, the most salient instance of coinsurance-capping regulation is the Medicare Part D market, where plans must adhere to a minimum standard benefit design that is characterized as a series of coinsurance rates. In 2018—the last year of our data—a standard Medicare Part D plan followed a four-phase structure with (i) a deductible, (ii) an initial coverage phase with a 25% coinsurance, (iii) a coverage gap phase with a 35% coinsurance, and (iv) a catastrophic coverage phase with a 5% coinsurance.[2] Plans also have the option of offering non-defined standard benefit designs (like tiered formularies) that differ from the structure above, but these offerings have to be actuarially equivalent (or more generous) than the standard design, and cost-sharing in each tier is subject to a "non-discriminatory" cap to prevent plans from discouraging high-cost beneficiaries from enrolling. Crucially, virtually no plans choose to pass rebates through to patients at the point of sale.[3] As a result, pa-

[1] An interesting exception, which we discuss in our model and results sections, is high-deductible health plans (HDHPs), whose design is constrained by a structural cap of 100% on coinsurance rates. Even though an insurance plan with a higher-than-100% coinsurance rate (even locally) appears perverse, instances of patients paying more for drugs through their insurance company while under the deductible than if they were to self-pay are commonly documented in the media.

[2] The coverage gap phase historically required patients to bear 100% of costs, but was slowly filled in starting with the passage of the Affordable Care Act in 2010, and culminating with the Bipartisan Budget Act of 2018. In 2025, provisions contained in the Inflation Reduction Act (IRA) removed the coverage gap phase from the Part D standard benefit design entirely. While these redesigns substantially lowered patient out-of-pocket costs—especially for sicker patients—they did not substantially alter the incentives identified in our theoretical model.

[3] In 2019, CVS Health offered a Medicare Part D plan—through its Silverscript Insurance Company—that promised to pass through a portion of rebates at point-of-sale. The plan was

tients enrolled in plans offering the same exact benefit design may face very different out-of-pocket costs if those plans cover drugs with different list prices (even if they have the same net-of-rebate price). As a concrete example, consider the case of two plans covering, respectively, Lantus and Basaglar: two equivalent, but differently priced versions of insulin glargine—a long-acting insulin widely used by patients with Type I and II diabetes. In 2018, the list price of Lantus was $400 for a pack of 5 pre-filled 3mL-syringes. The equivalent list price for Basaglar was $330. Hence, the out-of-pocket cost of a patient needing insulin glargine could be anywhere between $3.50 and $70 higher—depending on the coverage phase—for a patient enrolled in a plan covering Lantus but not Basaglar.[4]

To test the implications of our theoretical model, we study the relationship between changes in list price and changes in out-of-pocket costs across different types of plans, including standard Medicare Part D plans. We use prescription drug claims for a large set of commercially insured and Medicare patients. Our commercial data covers a range of plan types offered by self-insured employers, while our Medicare data includes a random sample of patients enrolled in both standalone Medicare Part D (PDP) and Medicare Advantage Part D (MA-PD) plans.[5]

The simple reduced-form relationship between list prices and out-of-pocket costs, however, may suffer from downward bias for two reasons. The first is endogeneity: drugs whose price increases faster may be more valuable. For example, a drug's price (both list and net-of-rebates) might spike after it obtains a new indication or after the publication of new evidence promoting its effectiveness. In those situations, we might expect plans to improve coverage for those drugs, dampening the relationship between cost and coverage. The second is that the payments that patients make under the deductible or through coinsurance do not depend directly on a drug's list price but on the reimbursement amount negotiated by the health plan (or PBM) with the pharmacy. This amount is related to a drug's list price, but also reflects a chain of margins and fees set along the distribution channel by wholesalers, pharmacies, and the PBM or plan.[6] As a result, changes in list price could be delayed and blunted by contracting

a commercial failure, only enrolling around 20,000 patients, and was withdrawn the following year (see, https://www.drugchannels.net/2019/01/new-part-d-enrollment-data-cvs-extends.html, accessed June 27th, 2026).

[4]As we discuss in detail below, patient cost-sharing is formally tied to the negotiated allowed amount rather than to the list price (WAC). Because the two are linked and the former varies across plans whereas the latter is set at the national level, we use list prices here for illustration.

[5]Both PDPs and MA-PDs cover exclusively retail prescription drug spending, but MA-PDs are attached to a privately administered health plan, while PDP enrollees are covered by Traditional Medicare.

[6]The payments on the distribution side are separate from the retrospective rebates that manufactur-

frictions along the supply chain, again leading to attenuation bias.

To address this bias, we use an instrumental variable research design that relies on exposure to the Medicaid segment as a plausibly exogenous shifter for list price growth. The Medicaid reimbursement formula includes an inflation penalty that strongly disincentivizes list price growth above inflation (Feng et al., 2023). Intuitively, our design compares the changes in out-of-pocket costs of drugs with high Medicaid sales—whose list price grows more slowly on average—to those of drugs with low Medicaid sales. The approach works under the exclusion restriction that other drug characteristics correlated with exposure to Medicaid are not driving the differential trend in out-of-pocket costs for commercial and Medicare patients.

Our estimates confirm that higher list prices lead to higher out-of-pocket costs, but, consistent with the theory, the effects are highly heterogeneous across plan types and market segments. We find the highest passthrough rates in Medicare Part D, where patients in plans that follow the standard benefit design or offer actuarially equivalent benefits experience a one-to-one passthrough rate—meaning that a 10 percent increase in list price translates to a 10 percent increase in out-of-pocket. Patients in enhanced plans experience a lower, but still significant passthrough rate of approximately 0.4, likely because many enhanced plans are still locally constrained by regulation in some of the phases.[7]

Conversely, we cannot detect any passthrough of list prices to out-of-pocket costs among commercial plans, with one exception. When we stratify the analysis by plan type, we see that patients in high-deductible health plans (HDHPs, covering around 10% of patients in the MarketScan data) also experience a one-to-one passthrough rate.[8] Like enhanced Medicare Part D plans, HDHPs also experience a local constraint, albeit not one generated by regulation. Rather, HDHPs are limited by the 100% coinsurance rate that is charged under the deductible. While it may seem odd that an insurance plan would optimally offer higher-than-100% coinsurance rates—essentially perverting the role of insurance—this is effectively what virtually all high-deductible plans do, as rebates for branded prescription drugs have reached, on average, 30% relative to list

ers pay to plans. Those rebates determine net-of-rebate prices but do not enter the point-of-sale allowed amount on which patient cost-sharing is based.

[7] In 2021, even though only 38% of Part D plans offered the standard benefit design, an additional 29% of plans charged the standard deductible amount (see, https://www.kff.org/medicare/medicare-part-d-a-first-look-at-medicare-prescription-drug-plans-in-2021/, accessed June 27th, 2026), and many enhanced plans also charged the maximum allowable coinsurance under the anti-discrimination provision (see, https://www.kff.org/medicare/key-facts-about-medicare-part-d-enrollment-premiums-and-cost-sharing-in-2021/, accessed June 27th, 2026).

[8] See Appendix C.2 for a breakdown of plan type market shares in the MarketScan data.

price (IQVIA Institute for Human Data Science, 2023).

Beyond passthrough, our model predicts that constrained plans should actively prefer to cover higher-list-price drugs, since doing so raises patient cost-sharing and helps relax the benefit-design constraint. We provide suggestive evidence consistent with this prediction by comparing how plans cover biologic drugs and their cheaper biosimilar equivalents. Using formulary data, we find that Medicare Part D plans are substantially more likely than commercial plans to grant preferential coverage to the higher-list-price biologics over their biosimilars—a gap of roughly 44 percentage points in a difference-in-difference comparison.

Our work complements recent empirical advances in the study of U.S. drug pricing, which tend to focus on net prices and treat list prices as exogenous (Olssen and Demirer, 2023; Feng and Maini, 2024; Ho and Lee, 2026).[9] Our main contribution is to provide a theory explaining why some U.S. health plans tie patient out-of-pocket costs for prescription drugs to an inflated price that does not reflect true acquisition cost. Existing theories of the role of list prices generally justify their existence by arguing that they act as an outside-option price for certain segments of the market. As a result, a high list price can increase profits if (i) some insured patients can still access the drug at those prices (Che, 2025), (ii) it boosts demand for health plans by worsening the outside option of going uninsured, and thus generating profits that both PBMs and manufacturers can share (Conti et al., 2026), and (iii) it supports price discrimination strategies.[10] While these theories are likely correct to some degree, however, none explains why list prices should be tied to out-of-pocket costs.

More broadly, our model fits within the literature studying how demand for insurance shapes benefit design, including a rich theoretical and empirical literature studying the effect of adverse selection on plan design (Rothschild and Stiglitz, 1976; Cutler and Reber, 1998; Glazer and McGuire, 2000; Finkelstein, 2004; Weyl and Veiga, 2017; Geruso et al., 2019; Shepard, 2022; Kreider et al., 2024), and a smaller literature studying how behavioral biases that lead patients to dislike premiums affect plan offerings (Abaluck and Gruber, 2016). Although no paper has specifically considered list prices as a tool in plan design, papers showing that Part D drug formularies are partially shaped by selection incentives are close conceptual neighbors to our work (Carey,

[9] An important exception is Che (2025), which studies list and net price setting in the market for anticoagulants but does not tackle the relationship between list prices and out-of-pocket costs.

[10] The lack of payer-specific rebate data means that no paper has studied price discrimination in the context of the US pharmaceutical market. However, evidence from similar markets exists (see, e.g., Grennan and Swanson, 2020).

2017; Lavetti and Simon, 2018; Geruso et al., 2019).

Finally, our results contribute to a growing literature on regulatory avoidance in health care markets. Much of this work studies how firms manipulate their behavior or reporting to extract larger payments: coding diagnoses more intensively in response to risk-adjustment (Geruso and Layton, 2020), timing patient discharges to reimbursement thresholds (Einav et al., 2018; Eliason et al., 2018), and gaming a subsidy's design (Decarolis, 2015). In contrast, our mechanism differs in that we show how plans can use a largely accounting variable to circumvent regulation that limits their strategy space. We are therefore closest to recent work on "tunneling," in which vertically integrated insurers inflate transfer prices to shift profits into reported drug costs and evade profit caps in Medicare Part D (Yde, 2025; Kakani et al., 2026), and nursing homes route profits to related parties to limit liability (Gandhi and Olenski, 2024). Where these firms tunnel profit into cost, plans in our setting tunnel cost onto patients, exploiting the list–net wedge to relax minimum-coverage rules while remaining nominally compliant.

2 A Model of Insurance Plan Design with List Prices and Rebates

We begin by introducing a model of insurance design where manufacturers and payers jointly determine list prices and out-of-pocket costs. We incorporate the standard incentives generated by adverse selection as in the seminal model by Rothschild and Stiglitz (1976), and we allow for the possibility that patients care more about premium dollars as shown by Abaluck and Gruber 2011 in the context of Medicare Part D plan choice. All the theoretical results we derive in this section hold whether or not we allow for this behavioral wedge—though they occasionally require slightly different regularity conditions. The results also hold without adverse selection (e.g., with perfect risk-adjustment) as long as the wedge is sufficiently large.

2.1 Setup

Our baseline framework adapts the binary-loss variation of the Rothschild and Stiglitz (1976) insurance model—as formalized in Hendren (2014)—to the setting of prescription drug insurance.

A unit measure of patients is indexed by a risk type θ distributed on $[0, 1]$ according to CDF $F(\theta)$ with continuous density $f(\theta) > 0$. The type θ is the probability that a

patient gets sick. If sick, the patient suffers a health shock of dollar value h, which can be fully offset by consuming one unit of a drug.

Patients can purchase insurance to subsidize the cost of the drug. Health insurance plans are characterized by a coinsurance rate α and a premium r. The coinsurance rate is applied to the list price of the drug, w, which is set by the drug manufacturer. Hence, a plan can be characterized as a pair $(\alpha w, r)$, where αw represents the actual out-of-pocket cost paid by patients who consume the drug. The drug manufacturer also sends a discount d to health plans, so the drug's net price to plans is $p = w - d$, leaving the net cost per fill for the health plan at $p - \alpha w$. We assume that the insurance market is perfectly competitive, so the premium is pinned down by the zero-profit condition

$$r = \mathbb{E}\left[\theta \middle| \text{enrolled}\right] \cdot (p - \alpha w) \tag{1}$$

We assume $h > p$, so there is no scope for moral hazard.[11]

Patient utility on a plan $(\alpha w, r)$ is

$$U(\theta, \alpha, w, r) = (1 - \theta) \cdot u(W - \kappa r) + \theta \cdot u(W - \kappa r - \alpha w),$$

where W is wealth, $u(\cdot)$ is increasing and concave, and $\kappa \geqslant 1$ is a behavioral parameter allowing for the possibility that patients put extra weight on premium dollars. Setting $\kappa = 1$ recovers the standard expected-utility benchmark. Uninsured patients can purchase the drug at an exogenously set price $\bar{p}$, so the utility of being uninsured is $U^0(\theta) = (1 - \theta) \cdot u(W) + \theta \cdot u(W - \bar{p})$.[12] We assume that the cash price faced by the uninsured weakly exceeds the net price paid by plans (i.e., $\bar{p} \geqslant p$), consistent with uninsured patients facing list-based prices at the pharmacy. Because—in equilibrium—effective cost-sharing never exceeds p in the analysis below, this assumption guarantees that insured patients never prefer to bypass their plan and pay the cash price, which simplifies some of our derivations, but has no meaningful implications on the results.

[11] In general, moral hazard would exacerbate the effects we describe in this model by amplifying the cost (and thus premium) differential between more and less generous insurance plans.

[12] We also assume that $h > \bar{p}$, so that uninsured patients also consume the drug. Treating $\bar{p}$ as exogenous shuts down the demand-side channel emphasized in Conti et al. (2026), in which higher list prices reduce the attractiveness of going uninsured and indirectly raise the demand for insurance. Empirically, manufacturers tend to neutralize this channel by expanding patient assistance programs as list prices rise. Treating $\bar{p}$ as exogenous lets us isolate the plan-side incentive without confounding it with the uninsured-demand channel, but does not otherwise affect the implications of the model.

2.2 When list prices are irrelevant

We first show that, in the absence of a regulatory cap on coinsurance, list prices are irrelevant. In other words, the payoff functions and strategies of the model are unchanged by perturbations of list prices that leave net prices intact.

To do so, we consider a perturbation from the baseline in which the drug's list price equals its net price ($w = p$, so the rebate is zero) where the manufacturer increases the list price to $w > p$ and pays a rebate $d = w - p$ to plans. This leaves the plan's acquisition cost unchanged. Uninsured patients continue to purchase the drug at price $\bar{p}$.

Proposition 1 ((Irrelevance of List Prices)). *Assume plans are restricted to contracts with positive premiums and effective cost-sharing that does not exceed the net price (i.e., $\alpha w \leqslant p$). Then any solution concept depending only on plans' strategy spaces and payoff functions yields identical predictions about $(\alpha w, r)$ outcomes across all list prices $w \geqslant p$.*

The proof of Proposition 1 relies on showing that for any potential contract $(\alpha w, r)$ and for any list price $w' \geqslant p$, α' can always be rescaled to reach $c = \alpha' w' = \alpha w$.

Two pillars underlie the Proposition: (i) list prices do not affect plan payoffs; and (ii) the set of feasible patient-relevant outcomes (c, r) does not depend on the list price, once attention is restricted to insurance contracts ($c \leqslant p$). The first pillar holds mechanically for any perturbation that holds the plan's net acquisition cost constant, as long as the uninsured price $\bar{p}$ is separately determined.

The second pillar fails when plan strategies are constrained, which can happen for two reasons. The first is if payers are limited by regulation, which we focus on in the next section. The second is if payers find it optimal to offer plans that exceed the boundaries assumed in the Proposition. While these boundaries generally appear relatively innocuous—they rule out plans where patients pay the insurance more than the insurance pays for drugs—they actually rule out a very common plan type in the U.S.: high-deductible health plans. These plans have "perverse" incentives, albeit only locally. Patients that pay for drugs under the deductible generally pay 100% of a drug's list price. Given widespread evidence that average rebates for branded drugs are around 30% (Aitken et al., 2016; Kakani et al., 2020), most of these patients are paying more money than it costs their insurance to acquire those same drugs. We do not engage with this specific case in the remainder of the section, because it is more complicated than our simple single-coinsurance model, but we return to it in the empirical analysis, when we isolate commercial HDHPs in our results.

2.3 When list prices matter: constrained benefit design

We now suppose that α is restricted to the interval $[0, \bar{\alpha}]$, with $\bar{\alpha} < 1$. We begin by analyzing a market with a single drug (drug A) and show that every Nash equilibrium of the free entry plan design game has all plans bunch at the cap $\bar{\alpha}$. For simplicity, we fix $w_A = p$, so that $p - \alpha w_A = (1-\alpha)\,p$. Setting $w_A > p$ would rescale the feasible effective cost-sharing cap to $\bar{\alpha} w_A$ and the zero-profit premium at the cap to $\mu\,(p - \bar{\alpha} w_A)$. When $\bar{\alpha} w_A < p$, so that effective cost-sharing stays below the net price, this reparametrization leaves the pooling-at-the-cap structure unchanged. We therefore lose no generality by normalizing $w_A = p$.

Lemma 2 ((Pooled Equilibrium under Caps)). *In the single-drug game (drug A with $w_A = p$), and under the assumption that gains to trade exist, together with mild regularity conditions on preferences and the type distribution, every Nash equilibrium has all active plans offering coinsurance $\alpha = \bar{\alpha}$ and a common premium $r^\star = \mu^\star \cdot p \cdot (1 - \bar{\alpha})$, where $\mu^\star = \mathbb{E}\,[\theta \mid \theta \geqslant \theta^\star]$ and $\theta^\star$ is the lowest type preferring this plan to remaining uninsured. If, in addition, the participation condition pinning down $\theta^\star$ admits a unique solution, the equilibrium is unique.*

Intuitively, plans are pushed towards the coinsurance cap by two forces. The first is the standard cream-skim argument of Rothschild and Stiglitz (1976), where any equilibrium with a pool at $\alpha^\star < \bar{\alpha}$ can be broken by a deviator offering $(\alpha^\star + \varepsilon, r^\star - \delta)$ that attracts the healthiest patients out of the original pool. The second is the premium wedge κ. Our formal proof in Appendix A.2 shows that either one of these two forces can generate the pooling result.[13]

Next, we consider a game where a second, identical drug is available (drug B) with list price $w_B > p$, but the same net price p as drug A (achieved by sending a rebate $d_B = w_B - p$ to plans). We also assume that both drugs are available to uninsured patients at the same exogenously set price $\bar{p}$.

Plans choose which drug $k \in \{A, B\}$ to cover in addition to their coinsurance $\alpha \in [0, \bar{\alpha}]$. The patient's effective cost-sharing on a plan covering drug k at coinsurance α is αw_k. While the cap restricts $\alpha \leqslant \bar{\alpha}$ uniformly across drugs, it does not restrict the product αw_k. For drug B, effective cost-sharing can reach $\bar{\alpha} w_B$, strictly above the max-

[13] In the absence of adverse selection, the premium wedge κ must be large enough to offset the risk aversion of patients near the participation margin—formally, a slope condition ($\kappa\mu^\star > 1$), which holds throughout the empirically relevant range $\kappa \in [2, 5]$. The pooling-at-the-cap conclusion fails without it; Appendix A.2 treats the knife-edge case $\kappa = 1$.

imum $\bar{\alpha}p$ available on drug A. Drug B therefore expands the set of feasible effective cost-sharing levels beyond what had been allowed by the cap in the single-drug game.

Proposition 3 ((List Prices as a Cap Workaround)). *In the two-drug game with* $w_B \in \left(p, \frac{p}{\bar{\alpha}}\right)$ *and under the same regularity conditions as Lemma 2, every Nash equilibrium has all active plans covering drug* B *at* $\alpha = \bar{\alpha}$*, with effective cost-sharing* $c^\star = \bar{\alpha} w_B$ *and premium* $r^{\star\star} = \mu^{\star\star}(p - \bar{\alpha} w_B)$*, where* $\mu^{\star\star} = \mathbb{E}[\theta | \theta \geqslant \theta^{\star\star}]$ *and* $\theta^{\star\star}$ *is the marginal enrollee under the new pool. If, in addition, the participation condition pinning down* $\theta^{\star\star}$ *admits a unique solution, the equilibrium is unique.*

Appendix A.3 provides a formal proof. The intuition behind the argument is that the two-drug game is isomorphic to the single drug game with the new effective cap at $\bar{\alpha} w_B$ instead of $\bar{\alpha} p$. This allows us to derive Proposition 3 as a corollary to Lemma 2.

Intuitively, Drug B allows health plans to circumvent the regulatory cap, which had foreclosed cream-skim deviations within drug A by preventing plans from raising α above $\bar{\alpha}$. The presence of a higher-list price drug provides an alternative path to the same cream-skim deviation. The increase in out-of-pocket costs relative to the single-drug benchmark is $c^\star - \bar{\alpha} p = \bar{\alpha}(w_B - p)$, which represents a one-for-one passthrough of the list-price markup, scaled by the regulatory cap.

2.4 Optimal list price setting

Proposition 3 treats w_B as an exogenously chosen list price above the net price p. We now show that the manufacturer of at least one drug in the two-drug game strictly prefers to set its list price above p, thus providing evidence that this scenario is relevant. We hold the net price p as constant throughout, and focus on modeling the list price-setting game between the manufacturers of drugs A and B.

Each manufacturer $k \in \{A, B\}$ chooses $w_k \geqslant p$ taking the rival's choice as given. Let $Q_k(w_A, w_B)$ denote the expected number of fills of drug k in the insurance-market equilibrium characterized in Lemma 2 and Proposition 3. Because each type-θ patient fills with probability θ, we have $Q_k = \mu_k \cdot M_k$, where M_k is enrollment and μ_k is the average risk of enrollees in drug-k plans. Manufacturer profits are

$$\pi_k(w_A, w_B) = p \cdot Q_k(w_A, w_B) - \Lambda(w_k - p) \tag{2}$$

where $\Lambda(\cdot)$ is an institutional penalty on the rebate capturing regulatory, reputational,

and negotiation frictions, with $\Lambda(0) = \Lambda'(0) = 0$.[14] We ignore the market for uninsured individuals.[15] At $w_A = w_B$ the two drugs are observationally equivalent (both deliver the same range of effective cost-sharing), so we assume $Q_A = Q_B$ in this case.

Proposition 4 ((Floor not Nash)). *Under the same regularity conditions as Lemma 2, the symmetric configuration $w_A = w_B = p$ is not a Nash equilibrium of the list-price-setting game. For any sufficiently small $\delta > 0$, either manufacturer can profitably deviate to $w_k = p + \delta$.*

The intuition follows from a perturbation argument. At $(w_A, w_B) = (p, p)$ the two drugs are observationally equivalent, and by the symmetric tie-breaking convention each manufacturer captures half of total expected fills $Q(p) = \mu^\star [1 - F(\theta^\star)]$ generated by the Lemma 2 pool. At $(p, p + \delta)$ with $\delta > 0$, Proposition 3 implies that all active plans cover drug B at the cap, so manufacturer B's expected fills jump discretely from $Q(p) \cdot \frac{1}{2}$ to $Q(p) + O(\delta)$. The resulting revenue gain is approximately $p \cdot Q(p) \cdot \frac{1}{2}$, whereas the costs are vanishingly small in δ. A symmetric argument applies to manufacturer A holding $w_B = p$.[16]

2.5 Discussion

The results of this model arise from the combination of: (i) incentives that push plans towards less generous plan offerings; and (ii) the presence of an exogenously set limit on benefit design, represented, in this case, by a coinsurance cap $\bar{\alpha}$. If coinsurance rates, or, more broadly, benefit design, could be freely chosen by insurers, our result would not apply, because increasing the list price would not expand the choice set of insurers. Therefore, our model suggests that list price increases are more likely to affect patients who are enrolled in plans that face generosity constraints, the most obvious of which arise from regulation.

Finally, we note how our simplified model with plans represented by one coinsurance rate applies to more complex real world plans. Unlike the linear cost sharing

[14] The presence of $\Lambda(\cdot)$ prevents list prices from going to infinity—a role similar to the political constraint parameter in Che (2025)—but otherwise does not affect Proposition 4.

[15] This is consistent with our previous treatment of $\bar{p}$ as exogenous. The uninsured revenue stream is likely to be second order and the existence of patient assistance programs decouples list prices and the prices faced by uninsured patients. Including profits from uninsured individuals does not alter the implications of the model, but introduces some additional mathematical complexity because increasing list price can affect the number of patients who are uninsured.

[16] While the Proof to Proposition 4 relies on a discontinuity in market shares, we derive the conditions under which a Nash equilibrium with list prices greater than net prices exists even in a world with smoother demand in Appendix A.5.

in our model, Medicare Part D plans are required to have a non-linear cost sharing structure, and many commercial plans also have such structures. A non-linear design means a plan can be constrained only "locally"—prevented from raising cost-sharing in one part of the design but not others—rather than globally. Our model does not address this case directly, but we expect its predictions to hold locally: plans may use list prices to circumvent the constraint in the specific parts of the design where it binds. We expect these local constraints to be particularly relevant in two settings. The first is among enhanced benefit design Part D plans, which, as mentioned previously, offer more generous benefits in some, but generally not all parts of the Part D benefit design. The second is high-deductible health plans, which are limited to a 100% coinsurance rate for the portion of spending under the deductible.[17]

3 Background, Data, and Empirical Approach

3.1 The role of list prices in the pharmaceutical supply chain

The term "list price" is informal, but generally refers to Wholesale Acquisition Cost (WAC). We follow this convention in our paper, and use the terms WAC and list price interchangeably throughout. WAC is a public manufacturer-posted price that aims to capture the price paid to drug manufacturers by wholesalers, which serves as a measurable starting point for the price paid on the distribution side of the industry. Moreover, as noted in both recent investigations by the U.S. Senate Finance Committee (Senate Finance Committee, 2021), and a 2024 Federal Trade Commission (FTC) public complaint against PBMs (Federal Trade Commission, 2024), contracts between PBMs and manufacturers include discounts expressed as a function of WAC.

However, while the idea that patients' out-of-pocket costs are tied to list prices is widely cited in the media and policy discussions, it is not, strictly speaking, correct. Instead, out-of-pocket costs are tied to a price negotiated by the payer (or PBM) and the dispensing pharmacy. This payment is usually referred to as "allowed amount" in claims data.

Figure 1 shows the steps that transform the WAC into the allowed amount, and the intermediaries involved.

Wholesalers, who pick up drugs from manufacturers and deliver them to pharma-

[17] Technically, all plans are constrained in such manner, but the large deductibles of high-deductible health plans mean that the constraint is much more likely to bind in a significant way.

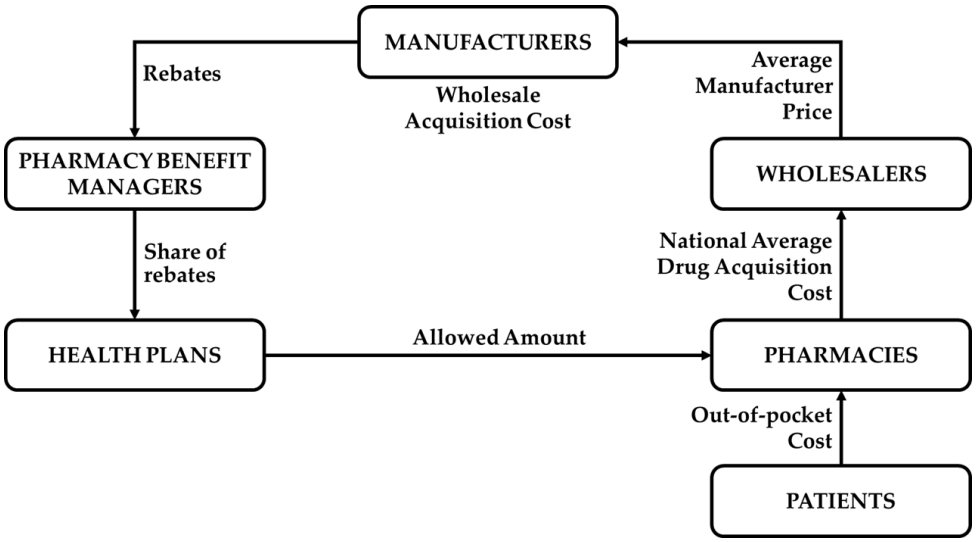

Figure 1: Overview of prices in the market for retail drugs

cies, are nominally charged the WAC, but typically negotiate small discounts. The average price paid by wholesalers is called the Average Manufacturer Price (AMP). AMP is reported to the government to comply with several programs (including the Medicaid Drug Rebate Program), but is not disclosed publicly (Peters, 2010). Government reports suggest that the difference between WAC and AMP is small for branded drugs, the focus of our study (between 4-10%), but larger for generics (around 40%) (Department of Health and Human Services, 2005; Congressional Budget Office, 2021).

Pharmacies then buy drugs from wholesalers at a markup above the wholesalers' acquisition cost. Although the acquisition costs for individual pharmacies are unknown, the *average* pharmacy acquisition cost is reported publicly as the National Average Drug Acquisition Cost (NADAC).[18]

Finally, pharmacies and PBMs (or plan sponsors) negotiate a reimbursement price that the pharmacy receives every time a patient fills a prescription. This reimbursement price is what appears under a prescription drug claim's allowed amount.

A patient's out-of-pocket cost can depend on the allowed amount if it is expressed as a coinsurance rate, or if the payment is made while the patient is under the deductible. Insofar as WAC affects the allowed amount, it can also affect the out-of-pocket cost. However, the extent of this influence depends on the contracts in place between manufacturers, wholesalers, pharmacies and payers, and on the timing of acquisitions at the successive steps in the distribution chains. We show in Appendix C.5 that the timing

[18]NADAC has been reported since 2013 by the Centers for Medicare and Medicaid Services and is downloadable from https://www.medicaid.gov/medicaid/nadac.

of NADAC and allowed amount increases often do not coincide. This may be because the payer or PBM contracts that set the allowed amount update only periodically, and because the inventory being dispensed was often acquired months earlier, so the cost basis underlying a given fill can itself lag the current posted price (both wholesalers and pharmacies must stock some amount of product).[19]

A separate, but related issue, is that the entities negotiating each contract may also have incentives to adjust margins for reasons that go beyond the value of the WAC. For example, Yde (2025) and Kakani et al. (2026) show that Medicare Part D plan sponsors that are vertically integrated with retail pharmacies increase the allowed amount on claims to "tunnel" profits from their insurer to their pharmacy arm—bypassing medical-loss-ratio regulation. In the same setting, Hofmann and Huang (2025) show that allowed amount on branded drug Part D claims increased following a 2010 reform that imposed a 50% manufacturer discount on certain claims. Finally, plans may also adjust their generosity in response to changes in WAC.

Hence, the relationship between what is commonly referred to as list price and patient cost is neither mechanical nor immediate. In the next section, we discuss how our empirical approach addresses these issues.

3.2 Empirical Approach

We estimate the impact of list price on various outcomes tied to patient out-of-pocket cost using the following regression design:

$$\Delta y_{jt} = \alpha_t + \gamma \times \Delta \ln (WAC)_{jt} + \varepsilon_{jt} \tag{3}$$

where j indexes drugs, and t indexes years. Δy_{jt} are changes in a list of dependent variables of interest, α_t are year fixed effects, and $\Delta \ln (WAC)_{jt}$ is the change in log WAC between periods $t-1$ and t. We use a first-difference specification because drug list prices can vary substantially and are usually reported in units that are not necessarily comparable across drugs, and because both our instrument and our object of interest concern the passthrough of year-over-year price changes.

There are two concerns in estimating this equation. First, as discussed in the background section, the effect of WAC on out-of-pocket cost is mediated through a series

[19]The lag between a NADAC increase and an allowed amount increase may also imply that wholesalers and pharmacies sometimes sell prescriptions at a loss, an issue which has been flagged in several media reports (see, e.g., https://www.drugchannels.net/2024/04/understanding-cvs-healthmckesson-and.html, accessed October 15th, 2025).

of contracts and lags that may dampen the short-run relationship between the two variables—even if in the long-run the passthrough were equal to 1. This dampening effect could bias the coefficient of the regression towards zero. We note that this bias is particularly problematic in regressions using variable *changes*, rather than levels, because changes generally occur on a one-off basis. Therefore accurately capturing the timing of changes is essential to identifying the correct effect.[20]

Second, list price is an equilibrium outcome, which makes a correlation with out-of-pocket cost difficult to interpret. For example, drugs might increase their list (and net) price in response to positive demand shocks (caused, e.g., by new clinical evidence, or label extensions). Such a shock is an omitted variable correlated with both price growth and coverage: if plans respond to higher drug value by improving the drug's coverage, a nominal price increase need not translate into higher out-of-pocket costs. The resulting omitted-variable bias is plausibly downward, though its precise sign depends on how the value shock affects allowed amounts, benefit design, and utilization.

Addressing these challenges requires an instrument that affects the incentive to increase or decrease list price, but is not otherwise correlated with other outcomes in the market.

Medicaid Market Share as an Instrument for List Price Growth

The instrument we propose is exposure to the Medicaid program, measured by Medicaid Market Share (MMS)—the share of a drug's sales that are paid for by Medicaid. Evidence from our previous work in Feng et al. (2023) suggests that list prices respond to incentives generated by the formula used by the Medicaid Drug Rebate Program (MDRP) to determine reimbursement for branded drugs.

The MDRP—which all manufacturers must abide by to participate in Medicaid—sets Medicaid's reimbursement for branded drugs as the AMP minus the sum of two rebates: a base rebate of 23.1% (or the highest discount granted to a commercial manufacturer) and an inflation rebate consisting of the difference between a drug's current AMP and its inflation-adjusted AMP at launch.[21] The formula implies that the base rebate is expressed as a function of the current AMP. As a result, faster-than-inflation

[20] For example, consider a situation where the list price of a drug increases on January 1st, but allowed amount does not update until April 1st. In that case, the average change in allowed amount relative to the prior year would be 25% lower than the average change in WAC.

[21] The statutory minimum rebate was equal to 15.1% up until December 2009, but was raised to 23.1% starting January 2010 by a provision in the Affordable Care Act. Discounts to Medicare Part D plans, the VA, and a few other institutional and public payers are exempt from the calculation of this rebate. See the Code of Federal Regulation (42 CFR § 447.509) for more details.

AMP growth *lowers* the overall Medicaid reimbursement.[22]

Because high AMP growth lowers Medicaid's reimbursement rates, drugs with a high Medicaid Market Share effectively face a penalty for increasing AMP, so they should, all else equal, exhibit lower AMP growth. Notice that because of this incentive, we should also expect these drugs to have a higher AMP at *launch* (Duggan and Scott Morton, 2006). Therefore, Medicaid exposure will work well as an instrument for changes in list price within drug, rather than level differences across drugs.[23]

While data on AMP is not publicly available, changes in WAC are closely tied to changes in AMP. Hence, our hypothesis is that changes in WAC will be negatively correlated with a drug's MMS.[24]

3.3 Data

To conduct the analysis outlined in the previous section, we combine data on drug prices and spending, and prescription drug claims. Data from all sources cover years 2007–2018 unless otherwise noted.

Drug spending and price data. We obtain drug list (WAC) and net prices for a set of approximately 1,000 brand drugs from SSR Health. SSR Health collects net sales on a quarterly basis from SEC filings of publicly traded firms and then calculates estimated US-level quarterly net prices by comparing net sales figures to invoice sales reported by a third-party (Symphony Health).[25] SSR Health data is at the product level, where a product is usually defined by a brand name (e.g., Lipitor).

We complement SSR Health with publicly available quarterly data on aggregated drug-level Medicaid spending from the Medicaid Public Use Files (PUF) published by the Centers for Medicare and Medicaid Services (CMS). The Medicaid PUF report national spending totals at the National Drug Code (NDC) level for all prescriptions filled

[22] In Appendix B.1, we reproduce Figure 2 of Feng et al. (2023), which illustrates graphically this somewhat counterintuitive implication of the formula.

[23] Another issue with this instrument is that, as our work in Feng et al. (2023) shows, manufacturers of drugs with a high MMS give lower rebates to commercial payers (though not Medicare Part D plans, which are excluded from the best-price clause). However, since this incentive does not change over time, it should not affect our specification, which focuses on price changes.

[24] Note that this result was already reported in Feng et al. (2023), albeit using a slightly different sample and specification.

[25] Not all drugs are available in all periods. Over time, as more firms have started reporting drug-specific earnings on their SEC filings, the sample has expanded. Occasionally, drugs are also dropped, mainly because of patent expiration. Also see Ippolito and Levy (2022) for a more detailed guide to the SSR Health data.

by Medicaid beneficiaries. These data can be matched to SSR Health data through a map provided by SSR Health. To calculate Medicare spending we use the Medicare Spending Dashboard, which reports overall drug-level spending levels in Medicare Part B and D.

Prescription drug claims data We use prescription drug claims from MarketScan, covering the commercial market, and from CMS, covering the Medicare market. MarketScan covers a large sample of patients enrolled in employer-sponsored insurance plans (over 47 million unique lives between 2007–2018). Each claim includes patient cost-sharing as well as information on the claim's total allowed amount. Cost-sharing distinguishes between payments made as copays, coinsurance, and under the deductible. The beneficiary information also includes the type of plan they are enrolled in (e.g., HMO, PPO, HDHP).

CMS data comes from the 20% random sample of Medicare Part D beneficiaries (over 90 million unique lives between 2007–2018). As with MarketScan, each claim also includes patient cost-sharing and total allowed amount. Cost-sharing is not split across different types of payments, but the standard benefit design of Medicare Part D means that the vast majority of payments are made as coinsurance or under the deductible. The main exception to this rule is for enhanced plans that provide more generous benefits, which occasionally cover some drugs using copays. We isolate these enhanced plans from basic plans in some of our analyses. When calculating out-of-pocket in Medicare claims, we exclude low-income subsidy (LIS) beneficiaries, who pay no, or limited cost-sharing.

Formulary coverage data

Data on formulary coverage come from MMIT Analytics, a firm that records drug tiering and restrictions from health-plan formularies. The data provide periodic snapshots of coverage at the drug–plan level across commercial, Medicare, Medicaid, and ACA-exchange plans, along with associated plan enrollment.[26] For each drug–plan pair we observe the formulary tier and any administrative restrictions (e.g., prior authorization or step therapy), which we use to measure whether a plan covers a drug. Our sample for this analysis uses a 2018 snapshot of biologic molecules that face biosimilar competition, together with their biosimilars.

[26] See Geruso et al. (2019) and Dean et al. (2024) for more detailed discussions of the MMIT data.

3.4 Sample selection

We create a core sample of branded, on-patent drugs based on four criteria.[27] First, since our analysis focuses on out-of-pocket costs patients face for drugs sold through the retail and mail channel, we exclude products primarily administered by physicians in outpatient settings. To implement this restriction we use Medicare Dashboard Spending data and select drugs whose Medicare sales are concentrated in Medicare Part D. Specifically, we exclude any product whose Medicare Part B sales exceed 20% of total Medicare sales.

Second, we exclude drugs whose invoice sales records are most likely to be underestimated. Undercounting of invoice sales can occur in the SSR Health data because the source of invoice sales consists of surveys that may not have perfect coverage of all sales channels. Undercounting is a particularly prominent issue for products that are sold through less common channels, such as specialty pharmacies. To avoid these issues we follow Kakani et al. (2020) and Feng et al. (2023) in screening out: (i) injectable drugs; (ii) drugs with an indication in oncology and (iii) drugs that record higher net sales than invoice sales.[28]

Third, we exclude contraceptives and smoking cessation drugs, because the ACA mandates that insurance plans cover them at no cost to patients. We confirm that our data reflect these rule changes.

Fourth, and finally, we restrict our sample to drugs that have been on the market for 15 years or less, applying the filter at the drug-year level. Older drugs can often accumulate inflation rebates so high that they exceed 100% of AMP. In those cases, the Medicaid program caps rebates at 100%, ensuring that manufacturers don't have to pay money to Medicaid on net, thus removing the incentive to lower list price growth.[29]

After applying these filters, we end up with a core sample of 304 drugs. Panel A of Table 1 provides drug-year level summary statistics for this sample, in order to provide context for our analysis. Drugs, on average, earn a little over eight percent of their invoice sales through Medicaid, albeit with significant variation across drugs. List price growth is about 12 percent per year. We note that WAC and allowed amount are often

[27] We only track drugs through the year before loss of exclusivity (LOE). After LOE, drug sales typically plummet, which means that list prices and cost-sharing are much less relevant.

[28] To be conservative, we filter out drugs only if their net sales are greater than invoice sales over the entire period in which a drug is in the SSR sample. We also exclude drugs whose net sales are more than 50 percent greater than its invoice sales in any given year other than the launch year and the year before loss of exclusivity.

[29] Appendix C.3 shows that our first stage results do not depend on this restriction, but that our instrument is strongest in the first 15 years after launch, thus justifying our approach.

Table 1: Summary Statistics

Panel A: Overall Statistics

Variables		Mean	SD	N
MMS		8.12%	10.73%	1940
WAC		106.58	574.63	1940
$\Delta \ln(\text{WAC})$		0.12	0.17	1636
Allowed Amount	Commercial plans	1,242.32	3,353.36	1921
	Medicare Plans	1,188.68	3,343.92	1728
$\Delta \ln(\text{Allowed Amount})$	Commercial plans	0.12	0.18	1618
	Medicare Plans	0.13	0.14	1443
Out of pocket cost	Commercial Plans	69.23	61.47	1921
	Medicare plans	163.80	256.29	1728
$\Delta \ln(\text{Out of pocket cost})$	Commercial Plans	0.08	0.20	1618
	Medicare plans	0.07	0.17	1443

Panel B: Statistics by Medicaid Exposure

		Below-median MMS			Above-median MMS		
Variables		Mean	SD	N	Mean	SD	N
MMS		1.84%	1.57%	923	13.82%	12.22%	1017
WAC		148.16	804.08	923	68.84	201.15	1017
$\Delta \ln(\text{WAC})$		0.14	0.23	771	0.10	0.08	865
$\Delta \ln(\text{Allowed Amount})$	Commercial plans	0.13	0.22	758	0.10	0.13	860
	Medicare Plans	0.15	0.17	643	0.11	0.10	800
$\Delta \ln(\text{Out of pocket cost})$	Commercial Plans	0.08	0.20	758	0.08	0.20	860
	Medicare plans	0.09	0.18	643	0.05	0.17	800

Notes: Summary statistics for the core sample, limited to drugs 15 years old or younger. Missing entries for out-of-pocket cost occur for drugs with low sales that do not appear in our claims data in a given year.

stated in different units (e.g., price per pill versus price per 30-day supply), which leads to the discrepancy in absolute values but not in growth rates. Our analysis focuses exclusively on growth rates within a drug, which removes fixed differences in units between WAC and claim-level measures.[30] Medicare and commercial plans have very similar allowed amounts, but Medicare plans—which are known to be less generous—have much higher average out-of-pocket costs.

Panel B provides summary statistics broken down by drugs' Medicaid Market Share

[30] We note that this approach does not account for time-varying changes in the composition of a drug's claims (e.g., shifts in average days supply, package or strength mix, or retail-versus-mail dispensing), which can move average claim-level amounts independently of list price.

(MMS). We find that list price growth is 4 percentage points higher in drugs with below median exposure to Medicaid relative to drugs with above median exposure. Growth in allowed amounts also has a similar quantitative difference. Finally, we find minimal differences across the two groups of drugs in average out-of-pocket spending growth for commercial plans, but a 4pp difference for Medicare plans.

4 Results

4.1 The effect of Medicaid Market Share on list price growth

As a first step, we confirm that MMS and WAC growth are inversely correlated. Formally, we run the following regression:

$$\Delta \ln (WAC_{jt}) = \delta_t + \beta \times MMS_{jt} + \varepsilon_{jt} \tag{4}$$

The ideal definition of MMS would be the fraction of a drug's *volume* sales that comes from Medicaid. SSR Health and CMS record units in different ways, making it difficult to compute the MMS directly. Instead, we compute MMS using invoice sales. As both Medicaid and invoice sales are reported at list price levels, the price component should cancel out when we take the ratio.[31] Formally, for any given drug j and period t, we define

$$MMS_{jt} = \frac{\text{Medicaid sales}_{jt}}{\text{Total invoice sales}_{jt}}$$

Table 2 shows estimates from the first-stage regression. Our findings confirm that drugs with a high MMS increase list price more slowly. Specifically, a standard deviation increase in MMS (approximately 10pp) results in lower list price growth by 1.6 pp each year.

F-statistic from our first-stage regression indicates a reasonably strong instrument. The inclusion of year fixed-effects, which we use in our main analysis, absorbs some of the variation in MMS, because the size of the Medicaid program has increased over time as a result of the gradual ACA expansion that began in 2012. Even so, the first-stage

[31] We estimate invoice sales from the SSR Health data by adjusting net sales by the list-to-net price ratio. In practice, we might still have some residual measurement error because Medicaid data incorporates fees to pharmacies. This may lead us to slightly overestimate MMS. We still prefer this approach to alternatives that have been used in other papers, such as relying on data from the Medical Expenditure Panel Survey (e.g., Ippolito and Levy, 2023), which, due to its small sample size, could introduce significant measurement error for less popular drugs.

Table 2: Evidence Related to the Instrument

Panel A – First Stage Estimates

	List Price Growth	
	(1)	(2)
MMS	-0.163	-0.158
	(0.033)	(0.033)
Year FE		X
N	1,636	1,636
R^2	0.011	0.047
Kleibergen-Paap F-stat	24.25	22.69

Notes: Estimates of Equation 4, using data from SSR Health on WAC to measure list price growth. Standard errors clustered at the drug level.

Kleibergen-Paap rk Wald statistic, which is robust to heteroskedasticity and within-drug clustering, remains above conventional thresholds (24.25 without and 22.69 with year fixed-effects). We nonetheless take two precautions against weak-instrument concerns. First, we estimate by limited-information maximum likelihood (LIML), which has better finite-sample bias properties than two-stage least squares under weak identification. Second, following Andrews et al. (2019), we report the Kleibergen-Paap F-statistic in all regression tables.[32]

4.2 Estimates of the Impact of List Price Growth on Out-of-Pocket Costs

To estimate the effect of list price growth on changes in out-of-pocket costs, we use MarketScan and CMS data from 2007–2018. We match claims to drugs based on the National Drug Code recorded on the claim. Both CMS and MarketScan report patient out-of-pocket in separate, clearly identifiable variables. Our main dependent variable is the log change in average claim-level out-of-pocket at the drug-year level. We also calculate separate measures of out-of-pocket cost by (i) plan type (PPO, HMO, and HDHP for MarketScan data, and basic vs. enhanced in Medicare data), and (ii) payment type (copay, coinsurance, and deductible—available only in the MarketScan data).

[32]The statistic varies across regressions because of variations in the sample of drugs used in each regression. The sample itself varies because of missing data in certain drugs, years and plan type or segment.

Throughout, we report results both by plan type and by payment type. The theory in Section 2 predicts passthrough wherever a patient's cost-sharing is tied to the allowed amount and the plan is constrained from lowering it—namely on the deductible and coinsurance margins—and no passthrough where cost-sharing takes the form of a fixed copay. Plan labels are therefore best understood as proxies for the mix of binding payment margins faced by a plan's enrollees, rather than as primitives: PPO and HMO enrollees pay disproportionately through copays, HDHP enrollees disproportionately under the deductible, and basic Part D enrollees almost entirely through coinsurance and the deductible. Where the data permit (the commercial sample), we therefore also estimate passthrough directly by payment type, which maps more tightly to the theoretical mechanism than the plan label does.

Two caveats apply to this mapping. First, the HDHP result operates through the under-deductible region, where the nominal 100% cost-sharing rule can push effective cost-sharing above the plan's net cost—a corner that lies outside the interior region of Proposition 3 (see Section 2.3). Second, enhanced Part D plans are constrained in some benefit phases but not others, so their intermediate estimate reflects a blend of constrained and unconstrained margins rather than a single regime.

The effect of list price growth on patient out-of-pocket costs in commercial plans

We start by focusing on commercial plans. Table 3 presents our results, which include regressions looking only at the claim-level out-of-pocket cost of patients in specific types of plans. In the OLS regressions, we find a significant but weak passthrough rate of WAC to out-of-pocket costs. Across all plans, we find that a 10% increase in list price is correlated with a 1.6% increase in out-of-pocket cost. We also find little variation across plan types.

Conversely, our instrumental variable esti

mates, despite being more imprecise, highlight large disparities in passthrough rates across plans. We find little evidence of passthrough in more generous PPO and HMO plans, but very high rates of passthrough (statistically indistinguishable from 1) in less generous HDHP plans.

A plausible explanation for these results is that patients in HDHP plans are more likely to make payments under the deductible and with coinsurance, which may be channels that apply more broadly. To unpack this further, we repeat the analysis, but study deductible, coinsurance, and copay claims separately.[33] Once again, the OLS es-

[33] 21.48% of deductible claims in our sample are partially paid by copay and 6.41% are partially paid

Table 3: Response in out-of-pocket Costs – Commercial Plans

	$\Delta \log(OOP)$													
	All Plans								PPO		HMO		HDHP	
	All claims		Deductible Claims		Coinsurance Claims		Copay Claims		All claims		All claims		All claims	
PANEL A. OLS ESTIMATES														
$\Delta \ln(WAC)$	0.156	0.164	0.262	0.239	0.370	0.376	0.141	0.134	0.174	0.181	0.0199	0.0366	0.175	0.197
	(0.077)	(0.076)	(0.1265)	(0.120)	(0.151)	(0.162)	(0.062)	(0.062)	(0.082)	(0.081)	(0.057)	(0.056)	(0.088)	(0.091)
PANEL B. INSTRUMENTAL VARIABLE ESTIMATES														
$\Delta \ln(WAC)$	-0.199	-0.107	0.709	0.567	1.078	0.999	-0.003	0.000	-0.203	-0.163	0.023	0.197	1.183	1.165
	(0.321)	(0.321)	(0.328)	(0.327)	(0.394)	(0.406)	(0.300)	(0.301)	(0.326)	(0.340)	(0.395)	(0.394)	(0.338)	(0.345)
Kleibergen-Paap F-stat	24.41	22.72	21.35	19.14	22.06	20.14	24.41	22.72	24.46	22.73	20.36	18.71	22.07	20.34
Year FE		X		X		X		X		X		X		X
N	1,618	1,618	1,569	1,569	1,585	1,585	1,618	1,618	1,616	1,616	1,588	1,588	1,586	1,586

Notes: Estimates of Equation 3 using claims data from MarketScan. Standard errors clustered at the drug level.

Table 4: Response in out-of-pocket Costs – Medicare Plans

	$\Delta \log(OOP)$					
	All Plans		Basic Plans		Enhanced Plans	
PANEL A. OLS ESTIMATES						
$\Delta \ln(WAC)$	0.280	0.244	0.277	0.236	0.228	0.187
	(0.108)	(0.119)	(0.088)	(0.100)	(0.112)	(0.120)
PANEL B. INSTRUMENTAL VARIABLE ESTIMATES						
$\Delta \ln(WAC)$	0.766	0.802	0.930	0.874	0.394	0.444
	(0.191)	(0.182)	(0.308)	(0.330)	(0.190)	(0.180)
Kleibergen-Paap F-stat	19.66	17.69	14.71	12.90	16.01	13.83
Year FE		X		X		X
N	1,443	1,443	1,379	1,379	1,414	1,414

Notes: Estimates of Equation 3 using CMS's 20% random sample of Medicare Part D beneficiaries. Basic plans follow the standard Medicare benefit design or offer actuarially equivalent benefits. Enhanced plans are more generous, and offer additional coverage. Standard errors clustered at the drug level.

timates are relatively low across all three payment types. Our IV estimates for copay claims show essentially no passthrough, whereas deductible and coinsurance claims have passthrough rates close to 1.

The effect of list price growth on patient out-of-pocket costs in Medicare plans

Next, we study Medicare plans. As previously discussed, all Medicare plans must provide a minimum basic coverage level or an actuarially equivalent design. Plans can also offer additional benefits, which usually amount to a lower deductible, additional coverage in the coverage gap, or lower coinsurance in the initial coverage phase. Similar to our commercial market analysis, we report average estimates across all plans and also report estimates split by coverage generosity.

Table 4 presents the results, which mimic the patterns seen in the commercial market. OLS estimates are relatively low and uniform across both basic and enhanced plans. Once again, we recover higher, though noisier, passthrough rates in our IV regression, which also shows higher passthrough for basic plans relative to enhanced plans.

by coinsurance. Claims are counted in a payment type bucket if any amount of the claim was paid with that payment type and therefore can be assigned to multiple payment categories.

Table 5: Changes in Allowed Amount

	Commercial		Medicare	
PANEL A. OLS ESTIMATES				
$\Delta \ln(WAC)$	0.552	0.536	0.486	0.478
	(0.218)	(0.213)	(0.209)	(0.208)
PANEL B. INSTRUMENTAL VARIABLE ESTIMATES				
$\Delta \ln(WAC)$	0.967	1.045	1.015	1.089
	(0.243)	(0.263)	(0.169)	(0.188)
Kleibergen-Paap F-stat	24.41	22.72	19.66	17.69
Year FE		X		X
N	1,618	1,618	1,443	1,443

Notes: Estimates of Equation 3, using allowed amount as the outcome. Commercial estimates are based on MarketScan claims and Medicare estimates are based on CMS's 20% random sample of Medicare Part D beneficiaries. Standard errors clustered at the drug level.

4.3 The relationship between list price and allowed amount

As shown above, we find very different estimates when using a standard panel regression versus an instrumental variables approach. We provide some direct evidence on contracting frictions, one of the two potential issues noted earlier.

To do so, we estimate the relationship between changes in list price and changes in the total allowed amount in a claim. As discussed in Section 2, the allowed amount is the negotiated price between the PBM (or payer) and the pharmacy that coinsurance and deductible payments are based on.

Formally, we estimate the specification in equation 3 using as dependent variable the log change in average claim-level allowed amount at the drug-year level. We run separate regressions for commercial and Medicare plans. Table 5 shows that the naive OLS estimates reflect a passthrough rate between 0.48 and 0.55 across the two samples. Conversely, our IV approach shows a much larger passthrough rate that is not statistically distinguishable from 1. This gap between the attenuated OLS estimate and the near-one IV estimate is consistent with contracting frictions and measurement lags dampening the short-run relationship.

4.4 Additional evidence on coverage of high-list-price drugs

To complement the evidence on the passthrough of list prices, we present suggestive evidence that Medicare Part D plans are more likely to cover high-list-price drugs than their low-list-price when both are present.

Our test compares coverage rates of biologic drugs and their biosimilar counterparts. Biologic drugs are large, structurally complex molecules produced in living cells and used to treat conditions such as autoimmune disorders, cancer, and diabetes. Biosimilars are products that are highly similar to, and have no clinically meaningful differences from, an already-approved reference biologic, and that can enter the market once the originator's period of exclusivity expires—playing a role analogous to that of generics for small-molecule drugs.

Empirically, biologics have had high list prices (and high rebates), while biosimilars have had lower list prices (and lower rebates). Despite this price advantage, biosimilars have generally been slow to gain market share. A leading explanation is that originator manufacturers leverage large rebates—made possible by high list prices—to secure preferred formulary placement and to discourage plans from covering lower-priced competitors, a practice often referred to as a rebate wall. This mechanism mirrors the one in our model: because rebates accrue to the plan while patient cost-sharing is tied to the list price, a plan can find it attractive to cover the high-list biologic even when the biosimilar is cheaper on a net-of-rebate basis.

We use formulary data from MMIT to compare coverage rates of biologics and their biosimilar counterparts across commercial and Medicare Part D plans. Our results (Figure2) show that across every single molecule, Part D plans are much less likely to cover biosimilars relative to biologics, whereas commercial plans essentially have parity coverage between the two (with the exception of insulin lispro).[34] A difference-in-difference estimate comparing the relative coverage rate of biologics and biosimilars across Part D and commercial plans suggests Part D plans are 44 percentage points more likely to give preferential coverage to the more expensive biologic product than commercial plans.

[34] Filgrastim, infliximab, interferons, and pegfilgrastim are more commonly administered by physicians; we include them because they are nonetheless listed on pharmacy formularies and are at times paid through the pharmacy benefit, for example under "white bagging" arrangements.

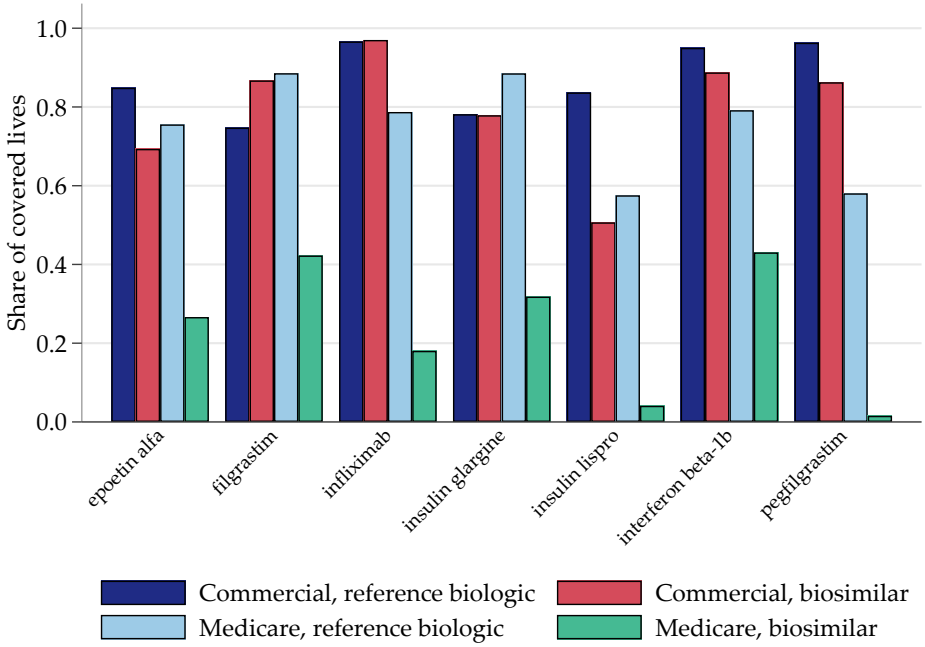

Figure 2: Formulary coverage of biologics and their biosimilars, by molecule and market.
Notes: *Each molecule group reports the share of lives enrolled in plans that cover the originator biologic and the share that cover its biosimilar(s), separately for commercial and Medicare Part D plans. Data refers to 2018.*

4.5 Discussion and Robustness Checks

Overall, these results are consistent with the predictions of our model. We find that list prices pass through to patient out-of-pocket costs, but only do so in a significant way for patients in plans that are constrained in their design: Medicare Part D plans—which must provide a minimum level of benefits, and commercial high-deductible health plans—which are constrained by the structural 100% coinsurance cap for payments under the deductible. Conversely, commercial PPO and HMO plans appear to leave cost-sharing unchanged.

Although the evidence is consistent with our model, there are alternative explanations for our findings. One could be that plans are simply being managed "passively," i.e., they do not change their deductible, copay, and coinsurance settings over time. Although passive management could happen in the commercial sample (where plans are picked by employers and drugs may represent a small fraction of overall spending), it seems unlikely that Medicare Part D plans, which are offered directly by insurers and specifically cover drugs, are not actively managed. Indeed, extensive evidence

shows that Part D plan sponsors carefully design their offerings along many margins, including (i) managing formularies in response to risk-adjustment (Carey, 2017; Lavetti and Simon, 2018) and manufacturer discounts in the coverage gap (Bertuzzi and Maini, 2024), (ii) adjusting their benefit design in response to the IRA (Cai et al., 2025), and (iii) strategically pricing to take advantage of a premium subsidy available to low-income beneficiaries (Decarolis, 2015).

Another possible explanation is that high- and low-MMS drugs display different trends in out-of-pocket costs for reasons other than their differential list price growth. While it is impossible to completely rule this out, we probe its importance by controlling for therapeutic-area fixed effects and for the change in the drug's net price. We include net price for two reasons: i) it is what actually affects plan costs; and ii) high-MMS drugs are also more exposed to the Medicaid best-price clause.[35] We caution, however, that net price is itself an equilibrium outcome. Conditioning on it can therefore introduce bad-control bias rather than cleanly isolating the exclusion concern, so we read this exercise as a robustness bound rather than a definitive adjustment. The estimates (reported in Appendix C.4) are quantitatively similar, albeit more imprecise. Adding therapeutic-area fixed effects in particular significantly weakens the first-stage F-stat, especially in the Medicare sample.

5 Conclusion

In this article, we provide both a theoretical justification for why and when payers pass through list price increases to patients in the form of higher out-of-pocket costs and empirical evidence consistent with the model. Our model highlights how adverse selection and premium overweighting can incentivize payers facing benefit design constraints to use higher list prices to raise patient cost-sharing. Using an instrumental-variables strategy based on incentives generated by Medicaid inflation rebates, we find that list prices pass through almost one-to-one to patient out-of-pocket spending in plans facing design constraints, but pass through much less in plans that can set cost-sharing freely.

Our findings have three main implications. First, our empirical results document that higher list prices translate into higher out-of-pocket costs, confirming that the list-

[35] As discussed in the background section, the best-price clause primarily affects net-price *levels*—high-MMS drugs tend to offer smaller commercial rebates, raising their net price—rather than net-price *growth*, so its first-order effect is differenced out of our specification. This concern also does not apply to Medicare Part D, whose discounts are exempt from the best-price calculation.

net price divergence is not a harmless accounting artifact (as sometimes argued by pharmacy benefit managers and plan sponsors) but a source of real financial burden, which could lower patient access to drugs. These findings support allegations in the FTC's recent complaint against large pharmacy-benefit managers (Federal Trade Commission, 2024), which argues that steering patients toward high-list-price products can cause harm even when these products cost the same (or less) because of higher rebates.

Second, however, our theory also suggests that plan-sponsor incentives alone are sufficient to sustain the list-to-net spread, independent of any contribution from PBMs. Our analysis adds some nuance to the claims in the FTC complaint. Although we do not model or rule out PBM incentives as a separate driver, our mechanism shows that the spread can persist on sponsor incentives alone. Our argument could help explain why the spread has persisted even as PBMs have increasingly integrated with insurers and moved towards full sharing of rebates with payers.[36]

Third, by bypassing benefit design regulation through the use of list prices, payers are effectively undoing some of the insurance protection guaranteed by the cap. In particular, because higher out-of-pocket costs allow plans to compete by lowering premiums, this workaround acts as a transfer from sick patients to healthy patients.[37]

Our findings also speak to the policy debate over how to contain rising list prices. Because list prices pass through to patient cost-sharing in constrained plans, measures that restrain list-price growth—such as the inflation rebates that now operate in both Medicaid and Medicare—should mechanically dampen the growth in out-of-pocket costs for the patients most exposed to this channel. However, two issues temper the potential gains to such policies. First, the restraining force of inflation rebates appears modest. Our estimates suggest that the incentives embedded in the Medicaid rebate formula slow annual list-price growth by only about 1.6 percentage points for a standard-deviation increase in Medicaid exposure, against a baseline of roughly 12 percent. Second, policies that cap prices can carry costs elsewhere in the supply chain, by limiting manufacturers' ability to react to shortages or surges in demand. More fundamentally, inflation rebates do not directly address the adverse selection and premium-overweighting incentives that lead constrained plans to exploit the list-net wedge in the first place. A more targeted remedy would sever the link between cost-sharing and list prices directly, either by requiring that rebates be passed through at the point of

[36] See, https://www.drugchannels.net/2023/08/surprising-data-on-employer-pbm-rebate.html, accessed June 27th, 2026.

[37] Interestingly, in plans with an out-of-pocket max, very sick patients are also unaffected by list price increases and can benefit from lower premiums.

sale, or by restricting cost-sharing to fixed copays.

References

Abaluck, Jason and Jonathan Gruber, "Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program," *American Economic Review*, June 2011, *101* (4), 1180–1210.

_ and _, "Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance," *American Economic Review*, August 2016, *106* (8), 2145–2184.

Aitken, Murray, Ernst R. Berndt, David Cutler, Michael Kleinrock, and Luca Maini, "Has The Era Of Slow Growth For Prescription Drug Spending Ended?," *Health Affairs*, September 2016, *35* (9), 1595–1603.

Andrews, Isaiah, James H. Stock, and Liyang Sun, "Weak Instruments in Instrumental Variables Regression: Theory and Practice," *Annual Review of Economics*, August 2019, *11* (Volume 11, 2019), 727–753.

Bertuzzi, Luca and Luca Maini, "Benefit Design And Biosimilar Coverage In Medicare Part D: Evidence And Implications From Recent Reforms," *Health Affairs*, May 2024, *43* (5), 717–724.

Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad, "What does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics*," *The Quarterly Journal of Economics*, August 2017, *132* (3), 1261–1318.

Cai, Christopher L., Anushka Bhaskar, Aaron S. Kesselheim, and Benjamin N. Rome, "Changes in Medicare Part D Plan Designs After the Inflation Reduction Act," *JAMA Internal Medicine*, October 2025, *185* (10), 1266–1273.

Carey, Colleen, "Technological Change and Risk Adjustment: Benefit Design Incentives in Medicare Part D," *American Economic Journal: Economic Policy*, February 2017, *9* (1), 38–73.

Chandra, Amitabh, Evan Flack, and Ziad Obermeyer, "The Health Costs of Cost Sharing," *The Quarterly Journal of Economics*, May 2024, *139* (4), 2037–2082.

Che, Catherine, "Do Middlemen Raise Drug Costs? Countervailing Market Power Meets Agency Frictions," 2025.

Congressional Budget Office, "A Comparison of Brand-Name Drug Prices Among Selected Federal Programs," February 2021.

Conti, Rena M., Brigham Frandsen, Michael Powell, and James B. Rebitzer, "Common Agent or Double Agent? Pharmacy Benefit Managers in the Prescription Drug Market," *The Review of Economics and Statistics*, May 2026, pp. 1–44.

Cutler, David M. and Sarah J. Reber, "Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection," *The Quarterly Journal of Economics*, May 1998, *113* (2), 433–466.

Dean, Emma B., Josh Feng, and Luca Maini, "Stocking Under the Influence: Spillovers from Commercial Drug Coverage to Medicare Utilization," July 2024.

Decarolis, Francesco, "Medicare Part D: Are Insurers Gaming the Low Income Subsidy Design?," *American Economic Review*, April 2015, *105* (4), 1547–1580.

Department of Health and Human Services, "Medicaid Drug Price Comparisons: Average Manufacturer Price to Published Prices," Technical Report June 2005.

Duggan, Mark and Fiona M. Scott Morton, "The Distortionary Effects of Government Procurement: Evidence from Medicaid Prescription Drug Purchasing," *Quarterly Journal of Economics*, 2006, *121* (1), 1–30.

Einav, Liran, Amy Finkelstein, and Neale Mahoney, "Provider Incentives and Healthcare Costs: Evidence From Long-Term Care Hospitals," *Econometrica*, 2018, *86* (6), 2161–2219. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA15022.

Eliason, Paul J., Paul L. E. Grieco, Ryan C. McDevitt, and James W. Roberts, "Strategic Patient Discharge: The Case of Long-Term Care Hospitals," *American Economic Review*, November 2018, *108* (11), 3232–3265.

Federal Trade Commission, "In the Matter of Caremark Rx, LLC; Zinc Health Services, LLC; Express Scripts, Inc.; Evernorth Health, Inc.; Medco Health Services, Inc.; Ascent Health Services LLC; OptumRx, Inc.; OptumRx Holdings, LLC; and Emisar Pharma Services LLC.," November 2024.

Feng, Josh and Luca Maini, "Demand Inertia and the Hidden Impact of Pharmacy Benefit Managers," *Management Science*, 2024, *70* (12), 8940–8961.

_ **, Thomas Hwang, and Luca Maini**, "Profiting from Most-Favored Customer Procurement Rules: Evidence from Medicaid," *American Economic Journal: Economic Policy*, 2023, *15* (2), 166–197.

Finkelstein, Amy, "Minimum standards, insurance regulation and adverse selection: evidence from the Medigap market," *Journal of Public Economics*, December 2004, *88* (12), 2515–2547.

Gandhi, Ashvin and Andrew Olenski, "Tunneling and Hidden Profits in Health Care," March 2024.

Geruso, Michael and Timothy Layton, "Upcoding: Evidence from Medicare on Squishy Risk Adjustment," *Journal of Political Economy*, March 2020, *128* (3), 984–1026.

_ **,** _ **, and Daniel Prinz**, "Screening in Contract Design: Evidence from the ACA Health Insurance Exchanges," *American Economic Journal: Economic Policy*, 2019, *11* (2), 64–107.

Glazer, Jacob and Thomas G. McGuire, "Optimal Risk Adjustment in Markets with Adverse Selection: An Application to Managed Care," *American Economic Review*, September 2000, *90* (4), 1055–1071.

Grennan, Matthew and Ashley Swanson, "Transparency and Negotiated Prices: The Value of Information in Hospital-Supplier Bargaining," *Journal of Political Economy*, April 2020, *128* (4), 1234–1268.

Hendren, Nathaniel, "Unravelling vs Unravelling: A Memo on Competitive Equilibriums and Trade in Insurance Markets," *The Geneva Risk and Insurance Review*, September 2014, *39* (2), 176–183.

Ho, Kate and Robin S. Lee, "Fisher-Schultz Lecture: Contracting Over Pharmaceutical Formularies and Rebates," *Econometrica*, 2026, *94* (3), 689–728. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA22488.

Hofmann, Katja and Zong Huang, "The Effect of Public Insurance Design on Pharmaceutical Prices: Evidence from Medicare Part D," *American Economic Journal: Economic Policy*, 2025, *forthcoming*.

Ippolito, Benedic and Joseph Levy, "Best Practices Using SSR Health Net Drug Pricing Data," *Health Affairs Forefront*, 2022.

Ippolito, Benedic N. and Joseph F. Levy, "The influence of Medicare Part D on the list pricing of brand drugs," *Health Services Research*, 2023, *58* (4), 948–952. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1475-6773.14139.

IQVIA Institute for Human Data Science, "The Use of Medicines in the U.S. 2023," Technical Report 2023.

Kakani, Pragya, Eric Yde, Genevieve Kanter, Richard G. Frank, and Amelia M. Bond, "Profit Regulation and Strategic Transfer Pricing by Vertically Integrated Firms: Evidence from Health Care," April 2026.

_ , Michael Chernew, and Amitabh Chandra, "Rebates in the Pharmaceutical Industry: Evidence from Medicines Sold in Retail Pharmacies in the U.S.," March 2020.

Kreider, Amanda R., Timothy J. Layton, Mark Shepard, and Jacob Wallace, "Adverse selection and network design under regulated plan prices: Evidence from Medicaid," *Journal of Health Economics*, September 2024, *97*, 102901.

Lavetti, Kurt and Kosali Simon, "Strategic Formulary Design in Medicare Part D Plans," *American Economic Journal: Economic Policy*, August 2018, *10* (3), 154–192.

Olssen, Alexander L. and Mert Demirer, "Drug Rebates and Formulary Design: Evidence from Statins in Medicare Part D," 2023.

Peters, Christie Provost, "Medicaid Payment for Generic Drugs," Issue Brief No. 839, National Health Policy Forum, Washington (DC) September 2010.

Roberts, Eric T., Jessica Phelan, Aaron L. Schwartz, Ellen Meara, Dominic Ruggiero, Lilly Estenson, Rachel M. Werner, and Jose F. Figueroa, "Loss of Subsidized Drug Coverage and Mortality among Medicare Beneficiaries," *New England Journal of Medicine*, May 2025, *392* (20), 2025–2034. _eprint: https://www.nejm.org/doi/pdf/10.1056/NEJMsa2414435.

Rothschild, Michael and Joseph Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *The Quarterly Journal of Economics*, 1976, *90* (4), 629–649.

Senate Finance Committee, "Insulin: Examining the Factors Driving the Rising Cost of a Century Old Drug," Technical Report 2021.

Shepard, Mark, "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange," *American Economic Review*, February 2022, *112* (2), 578–615.

Weyl, E. Glen and Andre Veiga, "Pricing Institutions and the Welfare Cost of Adverse Selection," *American Economic Journal: Microeconomics*, May 2017, *9* (2), 139–148.

Yde, Eric, "Vertical Integration and Regulated Profits in Pharmacy Benefits Markets," 2025.

A Proofs

A.1 Proof of Proposition 1

Proposition ((Irrelevance of List Prices)). *Assume plans are restricted to contracts with positive premiums and effective cost-sharing that does not exceed the net price (i.e.,* $c = \alpha w \leqslant p$*). Then any solution concept depending only on plans' strategy spaces and payoff functions yields identical predictions about* (c, r) *outcomes across all list prices* $w \geqslant p$*.*

Proof. This Proposition follows from a straightforward substitution. The patient's effective cost-sharing per fill is $c = \alpha w$. By construction, patient utility on a plan with strategy (α, r) and the plan's profit per enrollee both depend on w only through

$$(1-\theta)\cdot u\left(W-\kappa r\right)+\theta\cdot u\left(W-\kappa r-c\right)$$

$$\frac{\pi}{n}=r-\mu\left(p-c\right)$$

where $\mu = \mathbb{E}\left[\theta \middle| \text{enrolled in this plan}\right]$ is the average risk type of the plan's enrollees. Under the restriction $c \leqslant p$, the set of (c, r) pairs feasible to a plan is $[0, p] \times \mathbb{R}_+$ for every $w \geqslant p$, since $\alpha \in [0, 1]$ can always be scaled to reach any $c \leqslant p \leqslant w$. Hence the strategy space and payoff functions in (c, r) space coincide across all $w \geqslant p$, and any solution concept depending only on these objects yields identical predictions. □

A.2 Proof of Lemma 2

Lemma ((Pooled Equilibrium under the Cap)). *In the single-drug game (drug A with* $w_A = p$*), and under the assumption that gains to trade exist, together with mild regularity conditions on preferences and the type distribution, every Nash equilibrium has all active plans offering coinsurance* $\alpha = \bar{\alpha}$ *and a common premium* $r^\star = \mu^\star \cdot p \cdot (1-\bar{\alpha})$*, where* $\mu^\star = \mathbb{E}\left[\theta \middle| \theta \geqslant \theta^\star\right]$ *and* $\theta^\star$ *is the lowest type preferring this plan to remaining uninsured. If, in addition, the participation condition pinning down* $\theta^\star$ *admits a unique solution, the equilibrium is unique.*

Proof. Since $w = p$, effective cost-sharing at coinsurance α is $c = \alpha p$. Plans choose $\alpha \in [0, \bar{\alpha}]$ and charge a premium r. Free entry imposes a zero profit condition. A plan attracting the set of types $S \subseteq [0, 1]$ charges

$$r = \mathbb{E}\left[\theta \middle| \theta \in S\right] \cdot (1-\alpha)\, p. \tag{A1}$$

Patient utility on plan (α, r) is

$$U(\theta, \alpha, r) = (1-\theta)\, u(W - \kappa r) + \theta u(W - \kappa r - \alpha p),$$

and the outside option is $U^0(\theta) = (1-\theta)\, u(W) + \theta u(W - \bar{p})$.

We proceed in six steps.

1. **Single-crossing in α.** We establish the sorting property that underlies the cream-skim argument: higher-risk types have a stronger preference for lower coinsurance (more generous coverage), so that any deviation raising α attracts a "healthy slice" from the bottom of the pool.

   For any plan (α, r), define the absolute rate of substitution between coinsurance α and premium r as

   $$\mathrm{ARS}(\theta) = \frac{|\partial U/\partial \alpha|}{|\partial U/\partial r|} = \frac{\theta p u'(W - \kappa r - \alpha p)}{\kappa\left[(1-\theta)\, u'(W - \kappa r) + \theta u'(W - \kappa r - \alpha p)\right]},$$

   which is positive since both $\partial U/\partial \alpha < 0$ and $\partial U/\partial r < 0$. Write $a = u'(W - \kappa r)$ and $b = u'(W - \kappa r - \alpha p)$ with $b > a > 0$ (by strict concavity), so

   $$\mathrm{ARS}(\theta) = \frac{p}{\kappa} \cdot \frac{1}{1 + \frac{(1-\theta)a}{\theta b}}.$$

   The ratio $\frac{(1-\theta)a}{\theta b}$ is strictly decreasing in θ, so $\mathrm{ARS}(\theta)$ is strictly increasing in θ: sicker patients have a higher willingness to pay (in premium) for a marginal reduction in coinsurance.

2. **No equilibrium with a pool at $\alpha^\circ < \bar{\alpha}$.**

   Assume that the Nash equilibrium has all active plans pooling at some $\alpha^\circ \in [0, \bar{\alpha})$ with zero-profit premium $r^\circ = \mu^\circ \cdot (1 - \alpha^\circ)\, p$, where $\mu^\circ = \mathbb{E}[\theta \mid \theta \geqslant \theta^\circ]$ and θ° is the marginal type.

   Choose $\varepsilon > 0$ small enough that $\alpha^\circ + \varepsilon \leqslant \bar{\alpha}$. A deviating plan offers $(\alpha^\circ + \varepsilon, r_{\mathrm{dev}})$, where r_{dev} is set so that a type $\hat{\theta} \in (\theta^\circ, 1)$ is indifferent between the deviating plan and the incumbent pool:

   $$U\left(\hat{\theta}, \alpha^\circ + \varepsilon, r_{\mathrm{dev}}\right) = U\left(\hat{\theta}, \alpha^\circ, r^\circ\right).$$

   By single crossing (Step 1), every type $\theta < \hat{\theta}$ in the original pool strictly prefers the

deviating plan. The deviator attracts at least the types in $\left[\theta^\circ, \hat{\theta}\right)$ from the incumbent pool. Some types below θ° who were previously uninsured may also find the deviation attractive, since it offers a lower premium than the plan they rejected (despite higher coinsurance). If any such types enroll, they have $\theta < \theta^\circ < \mu^\circ$, so they further lower the deviator's average risk and strengthen profitability. This set has positive measure since $f(\theta) > 0$. The mean risk of these types is

$$\mu_{\text{dev}} = \mathbb{E}\left[\theta | \, \theta \in \left[\theta^\circ, \hat{\theta}\right)\right] < \mu^\circ,$$

strictly below the pool mean. The deviator's zero-profit premium is $r^0_{\text{dev}} = \mu_{\text{dev}} \cdot (1 - \alpha^\circ - \varepsilon)\, p$, and the premium it can charge (pinned by the indifference condition) satisfies $r_{\text{dev}} > r^0_{\text{dev}}$ for ε sufficiently small. Because the cost advantage $\mu^\circ - \mu_{\text{dev}} > 0$ is bounded away from zero while the premium concession is $O(\varepsilon)$, the deviator earns strictly positive profit, contradicting the assumption that the pool at α° is a Nash equilibrium.

3. **The pool at $\alpha = \bar{\alpha}$ cannot be cream-skimmed.**

   We show that the candidate equilibrium admits no profitable deviation. The regulatory cap rules out an upwards deviation in the coinsurance *($\alpha > \bar{\alpha}$)*. A deviation that keeps the coinsurance but alters the premium is also not relevant. A plan offering $(\bar{\alpha}, r_d)$ with $r_d < r^\star$ attracts the full pool but earns a strictly negative profit. A plan with $r_d > r^\star$ attracts no one.

   Therefore, we focus on downward deviations that attract sicker types ($\alpha_d < \bar{\alpha}$). A plan offering (α_d, r_d) with $\alpha_d < \bar{\alpha}$ provides more generous coverage at a higher premium. By single crossing, a type $\tilde{\theta}$ indifferent between the deviation and the incumbent pool satisfies

$$U\left(\tilde{\theta}, \alpha_d, r_d\right) = U\left(\tilde{\theta}, \bar{\alpha}, r^\star\right),$$

   and types above $\tilde{\theta}$ (sicker) strictly prefer the deviation. The deviator therefore attracts types $\theta \in \left(\tilde{\theta}, 1\right]$ with mean risk

$$\mu_d = \mathbb{E}\left[\theta | \, \theta \in \left(\tilde{\theta}, 1\right]\right] > \mu^\star.$$

   The zero-profit premium for this pool is $r^0_d = \mu_d \cdot (1 - \alpha_d)\, p$. For the deviation to earn positive profit, the premium must satisfy $r_d > r^0_d$.

We show this fails by comparing slopes in (c, r) space, where $c = \alpha p$ is effective cost-sharing. The zero-profit locus for the pool $\left[\tilde{\theta}, 1\right]$ is the line $r = \mu_d (p - c)$ with slope $-\mu_d$. The indifference curve of type $\tilde{\theta}$ through the incumbent point $(\bar{\alpha}p, r^\star)$ has slope $-\text{ARS}\left(\tilde{\theta}\right)/p$ (from Step 1), and since $\text{ARS}\left(\tilde{\theta}\right) < p/\kappa$ for all $\tilde{\theta} < 1$ at every (c, r), the indifference curve is shallower than the zero-profit line whenever $1/\kappa < \mu_d$. Since $\mathbb{E}\left[\theta | \theta \geqslant \tilde{\theta}\right]$ is increasing in $\tilde{\theta}$ and $\tilde{\theta} \geqslant \theta^\star$, the condition $\kappa\mu^\star > 1$ is sufficient.[38]

The slope bound from Step 1 holds pointwise: $\text{ARS}(\theta) < p/\kappa$ at every (c, r), not only at the incumbent point, since the factor $\frac{1}{1+((1-\theta)a/\theta b)}$ is strictly below one for every $\theta < 1$. Hence along type $\tilde{\theta}$'s indifference curve through $(\bar{\alpha}p, r^\star)$ the slope satisfies

$$\left|\frac{dr}{dc}\right| = \text{ARS}\left(\tilde{\theta}\right)/p < \frac{1}{\kappa} < \mu_d$$

at every $c < \bar{\alpha}p$, where the last inequality uses $\kappa\mu_d \geqslant \kappa\mu^\star > 1$. At $c = \bar{\alpha}p$ the indifference curve passes through $r^\star = \mu^\star(p - \bar{\alpha}p)$, which lies strictly below the deviator's zero-profit line $r = \mu_d(p - c) = \mu_d(1 - \bar{\alpha})p$ because $\mu_d > \mu^\star$; and since the line rises strictly faster than the indifference curve as c falls, the curve remains strictly below it for all $c < \bar{\alpha}p$. Therefore the indifference premium r_d is strictly below the break-even premium r_d^0 for any deviation $\alpha_d < \bar{\alpha}$, and the deviator cannot earn positive profit.

Remark on $\kappa = 1$. At $\kappa = 1$, the condition $\kappa\mu^\star > 1$ reduces to $\mu^\star > 1$, which fails since $\mu^\star = \mathbb{E}[\theta | \theta \geqslant \theta^\star] < 1$. The slope comparison must then hold directly. Case (b) deviations are unprofitable if and only if

$$\frac{\text{ARS}\left(\tilde{\theta}\right)}{p} < \mathbb{E}\left[\theta | \theta \geqslant \tilde{\theta}\right] \qquad \text{for all } \tilde{\theta} \in [\theta^\star, 1), \tag{R1}$$

a joint restriction on preferences *and* the type distribution. At $\kappa = 1$, $\text{ARS}\left(\tilde{\theta}\right)/p > \tilde{\theta}$ always holds (since $b > a$), so R1 requires the conditional mean to exceed $\tilde{\theta}$ by enough, and it can fail when F concentrates mass just above an interior point, and a deviator targeting that slice breaks the pool.

The limit $\tilde{\theta} \to 1$ (a deviation targeting a thin slice of enrollees) yields a necessary

[38] The condition $\kappa\mu^\star > 1$ is satisfied throughout the empirically relevant range $\kappa \in [2, 5]$ for any type distribution with $\mathbb{E}[\theta | \theta \geqslant \theta^\star] > 1/\kappa$. At $\kappa = 1$, the bound $\text{ARS}\left(\tilde{\theta}\right) < p$ does not suffice on its own; the argument requires a joint restriction on $u(\cdot)$ and F (condition R1), stated in the Remark below.

condition. Writing $b/a = u'(W - r^\star - \bar{\alpha}p)/u'(W - r^\star)$ for the marginal utility ratio and expanding to first order in $\varepsilon = 1 - \tilde{\theta}$ gives $\text{ARS}/p \approx 1 - \varepsilon(a/b)$ while $\mu_d \approx 1 - \varepsilon/2$ (for any distribution with continuous, positive density at $\theta = 1$), so the slope condition holds in the limit whenever $a/b > 1/2$, i.e., $b/a < 2$.

For the uniform distribution, this necessary condition is also sufficient: at $b/a = 2$, R1 reduces to $2\tilde{\theta}/\left(1+\tilde{\theta}\right) < \left(1+\tilde{\theta}\right)/2$, i.e., $\left(1-\tilde{\theta}\right)^2 > 0$, which holds for all $\tilde{\theta} < 1$ with the limit $\tilde{\theta} \to 1$ as the unique binding case (and is slack everywhere for $b/a < 2$). For CRRA utility with coefficient γ, $b/a = (1 - \bar{\alpha}p/(W - r^\star))^{-\gamma}$, so the condition $b/a < 2$ is equivalent to $\gamma < \frac{\ln 2}{\ln\left(\frac{1}{1-\bar{\alpha}p/(W-r^\star)}\right)}$. At a drug-cost-to-income ratio $\bar{\alpha}p/(W - r^\star) = 0.03$, this allows $\gamma \leqslant 22$; at 0.10, it allows $\gamma \leqslant 6.6$. At $\kappa = 1$ we therefore maintain R1, which is satisfied, e.g., by CRRA preferences with empirically standard risk-aversion coefficients ($\gamma \leqslant 5$) together with a uniform (or near-uniform) type distribution.

4. **No separating equilibrium.**

   Suppose a Nash equilibrium features multiple active contracts $\{(\alpha_j, r_j)\}_{j=1}^J$ with $\alpha_1 < \cdots < \alpha_J \leqslant \bar{\alpha}$ and $J \geqslant 2$. By single crossing, each contract attracts a connected interval of types, and less generous contracts (higher α_j) attract healthier (lower-θ) types. The type intervals are therefore ordered opposite to the coinsurance index: contract J (least generous) occupies the healthiest interval at the bottom, and contract 1 (most generous) the sickest types up to $\theta = 1$. Write contract j's interval as $\left[\underline{\theta}_j, \bar{\theta}_j\right]$, with $\underline{\theta}_J \leqslant \bar{\theta}_J = \underline{\theta}_{J-1} \leqslant \cdots \leqslant \bar{\theta}_1 = 1$.

   Consider the least generous sub-cap contract—the active contract with the highest $\alpha_j < \bar{\alpha}$ (it exists because $\alpha_1 < \cdots < \alpha_J \leqslant \bar{\alpha}$ with $J \geqslant 2$ leaves at most one contract at the cap). Apply the Step 2 cream-skim to its pool: a deviator offering $(\alpha_j + \varepsilon, r_{\text{dev}})$, with ε small enough that $\alpha_j + \varepsilon \leqslant \bar{\alpha}$, attracts the healthiest types of $\left[\underline{\theta}_j, \bar{\theta}_j\right]$ at strictly positive profit, breaking it. Hence no active contract can sit strictly below the cap, and all active plans pool at $\bar{\alpha}$ (which, by Step 3, is itself not cream-skimmable), collapsing the menu to a single contract.

5. **Existence.**

   Steps 1–3 establish that any Nash equilibrium must pool at $\bar{\alpha}$. We verify the candidate equilibrium is internally consistent.

   (a) *Participation.* Define $V(\theta) \equiv U(\theta, \bar{\alpha}, r^\star) - U^0(\theta)$. At $\theta = 0$: $V(0) = u(W - \kappa r^\star) -$

$u(W) < 0$. The gains-to-trade assumption implies $V(\theta) > 0$ for θ near 1. By the intermediate value theorem, $\theta^\star \in (0,1)$ exists.

(b) *Monotonicity.* Since $U(\theta, \bar{\alpha}, r^\star) = u(W - \kappa r^\star) + \theta\left[u(W - \kappa r^\star - \bar{\alpha} p) - u(W - \kappa r^\star)\right]$ and $U^0(\theta) = u(W) + \theta\left[u(W - \bar{p}) - u(W)\right]$, $V(\cdot)$ is affine in θ. Combined with $V(\theta^\star) = 0$ and $V(1) > 0$, this gives $V' > 0$. All $\theta > \theta^\star$ strictly prefer the pool to the outside option.

(c) *Zero profit.* By construction, $r^\star = \mu^\star(1 - \bar{\alpha})\, p$.

(d) *Fixed-point consistency.* Define the mapping $\theta \mapsto T(\theta)$ as the solution to $U(T, \bar{\alpha}, \mathbb{E}[\theta' | \theta' \geqslant \theta](1 - \bar{\alpha})\, p) = U^0(T)$. The map T is continuous in θ (U, U^0, and the pooled premium $\mathbb{E}[\theta' | \theta' \geqslant \theta](1 - \bar{\alpha})\, p$ all vary continuously, the last because $f > 0$). At $\theta = 0$ the pooled premium is strictly positive while a type-0 enrollee never files a claim, so $U(0, \bar{\alpha}, \cdot) = u(W - \kappa r) < u(W) = U^0(0)$ and the indifferent type is strictly positive: $T(0) > 0$. At $\theta = 1$ the premium is $(1 - \bar{\alpha})\, p$, and the gains-from-trade assumption makes the sickest type strictly prefer the pooled contract to remaining uninsured, so $T(1) < 1$. Writing $g(\theta) \equiv T(\theta) - \theta$, continuity together with $g(0) > 0$ and $g(1) < 0$ yields, by the intermediate value theorem, a fixed point $\theta^\star = T(\theta^\star)$.

6. **Uniqueness of the participation margin.**

Steps 1–3 imply that any Nash equilibrium pools at $\bar{\alpha}$, so equilibria can differ only in the participation margin, which must solve the fixed-point equation

$$\Psi(\theta) \equiv U(\theta, \bar{\alpha}, \mathbb{E}[\theta' | \theta' \geqslant \theta](1 - \bar{\alpha})\, p) - U^0(\theta) = 0.$$

Multiple solutions cannot be ruled out in general. Suppose $\theta_0 < \theta_1$ are both solutions, with pools $[\theta_0, 1]$ and $[\theta_1, 1]$, mean risks $\mu_0 < \mu_1$, and zero-profit premiums $r_0 = \mu_0(1 - \bar{\alpha})\, p < r_1 = \mu_1(1 - \bar{\alpha})\, p$. Against the candidate equilibrium at θ_1, an entrant offering $(\bar{\alpha}, r_0 - \varepsilon)$ attracts the pool $[\theta'(\varepsilon), 1]$ with $\theta'(\varepsilon)$ slightly below θ_0, and earns profit per enrollee

$$r_0 - \varepsilon - \mathbb{E}[\theta | \theta \geqslant \theta'(\varepsilon)](1 - \bar{\alpha})\, p = (\mu_0 - \mathbb{E}[\theta | \theta \geqslant \theta'(\varepsilon)])(1 - \bar{\alpha})\, p - \varepsilon,$$

in which both terms are $O(\varepsilon)$. Whether entry is profitable depends on whether the average-cost schedule is flatter or steeper than the demand schedule at the crossing, the typical source of multiple equilibria in selection markets. We therefore impose, as a maintained regularity condition, that Ψ admits a unique root on

$(0, 1)$, equivalently that the demand and average-cost schedules cross only once. Under this condition, Steps 1–4 deliver a unique equilibrium. We emphasize that the coverage prediction (pooling at $\alpha = \bar{\alpha}$) holds in *every* equilibrium regardless of this condition. Only the participation margin requires it.

□

A.3 Proof of Proposition 3

Proposition ((List Prices as a Cap Workaround)). *In the two-drug game with $w_B \in \left(p, \frac{p}{\bar{\alpha}}\right)$ and under the same regularity conditions as Lemma 2, every Nash equilibrium has all active plans covering drug B at $\alpha = \bar{\alpha}$, with effective cost-sharing $c^\star = \bar{\alpha} w_B$ and premium $r^{\star\star} = \mu^{\star\star}(p - \bar{\alpha} w_B)$, where $\mu^{\star\star} = \mathbb{E}[\theta | \theta \geqslant \theta^{\star\star}]$ and $\theta^{\star\star}$ is the marginal enrollee under the new pool. If, in addition, the participation condition pinning down $\theta^{\star\star}$ admits a unique solution, the equilibrium is unique.*

Proof. We reparametrize the two-drug game in terms of effective cost-sharing $c = \alpha w_k$. As established in the proof of Proposition 1, patient utility and plan profit depend on a plan's strategy (k, α, r) only through (c, r), because the rebate $d_k = w_k - p$ offsets the higher pharmacy payment (net cost per fill is $p - c$ regardless of which drug is covered).

The set of c-values feasible to a plan covering drug k at coinsurance $\alpha \in [0, \bar{\alpha}]$ is $[0, \bar{\alpha} w_k]$. Across both drugs, the feasible c-set is therefore $[0, \bar{\alpha} w_A] \cup [0, \bar{\alpha} w_B] = [0, \bar{\alpha} w_B]$, since $w_B > w_A = p$.

For $w_B \in \left(p, \frac{p}{\bar{\alpha}}\right)$, we have $\bar{\alpha} w_B < p$, so the upper bound $c \leqslant \bar{\alpha} w_B$ binds before the zero-profit constraint $c \leqslant p$ (which requires $r = \mu(p - c) \geqslant 0$). The game in (c, r) space is therefore isomorphic to a single-drug game with an effective cap at $\bar{\alpha} w_B$: the feasible set is $[0, \bar{\alpha} w_B] \times \mathbb{R}_+$ and the payoff functions are identical to the single-drug setup with cap $\bar{\alpha} w_B / p$ applied to the net price p.

By Lemma 2, every equilibrium of this reparametrized game pools at the effective cap, $c^\star = \bar{\alpha} w_B$. This pooling level requires $\alpha = \bar{\alpha} w_B / p > \bar{\alpha}$ on drug A (since $w_B > p$), which is infeasible. It is achievable only by covering drug B at $\alpha = \bar{\alpha}$.

The equilibrium premium is determined by the zero-profit condition:

$$r^{\star\star} = \mu^{\star\star}(p - \bar{\alpha} w_B),$$

where $\mu^{\star\star} = \mathbb{E}[\theta | \theta \geqslant \theta^{\star\star}]$ and $\theta^{\star\star}$ is the marginal type indifferent between the pool and the outside option. Existence of $\theta^{\star\star}$ follows from the fixed-point argument in Step 5 of

the proof of Lemma 2. Uniqueness holds under the maintained participation-margin regularity condition (Step 6), applied to the reparametrized game. □

A.4 Proof of Proposition 4

Proposition ((Floor not Nash)). *Under the same regularity conditions as Lemma 2, the symmetric configuration $w_A = w_B = p$ is not a Nash equilibrium of the list-price-setting game. For any sufficiently small $\delta > 0$, either manufacturer can profitably deviate to $w_k = p + \delta$.*

Proof. At $(w_A, w_B) = (p, p)$, the two drugs are observationally equivalent and the insurance-market equilibrium is the Lemma 2 pool at $\bar{\alpha}$ with total expected fills $Q(p) = \mu^\star [1 - F(\theta^\star)]$. By the symmetric tie-breaking convention, each manufacturer captures $Q_k = Q(p)/2$.

Consider a deviation by manufacturer B to $w_B = p + \delta$ for $\delta > 0$. For δ small enough that $p + \delta < p/\bar{\alpha}$, Proposition 3 applies: the unique insurance-market equilibrium has all active plans covering drug B at $\alpha = \bar{\alpha}$.

Revenue gain. Manufacturer B's expected fills jump from $Q(p)/2$ to $Q_B(p, p + \delta)$. Since the new pool covers all insured types (with a marginal type $\theta^{\star\star}$ satisfying $\theta^{\star\star} \to \theta^\star$ as $\delta \to 0$ by continuity of the equilibrium), we have $Q_B(p, p + \delta) = Q(p) + O(\delta)$, where the $O(\delta)$ term reflects the movement of the participation margin; its sign depends on whether the premium reduction or the higher cost-sharing dominates for the marginal type, and is immaterial to the argument. The revenue gain is

$$\Delta\pi_B^{\text{rev}} = p \cdot [Q(p) + O(\delta)] - p \cdot \frac{Q(p)}{2} = \frac{p \cdot Q(p)}{2} + O(\delta),$$

which is bounded away from zero—an $O(1)$ gain from capturing the rival's half of the market.

Penalty cost. The institutional penalty is $\Lambda(\delta)$. By assumption, $\Lambda(0) = \Lambda'(0) = 0$, so Taylor expansion gives $\Lambda(\delta) = O(\delta^2) = o(\delta)$.

Net gain. The change in manufacturer B's profit is

$$\Delta\pi_B = \frac{p \cdot Q(p)}{2} + O(\delta) - o(\delta) > 0$$

for all $\delta > 0$ sufficiently small. A symmetric argument applies to manufacturer A deviating to $w_A = p + \delta$ while holding $w_B = p$. Therefore, (p, p) is not a Nash equilibrium. □

A.5 List Price Setting Under Smoother Demand Functions

To move to a less stark setup, we assume that $Q_k(w_A, w_B)$ varies smoothly in list prices.

Define $\Phi(w_A, w_B) = p \cdot \frac{\partial Q_B}{\partial w_B}$ as the marginal revenue gain from a small list-price increase (Φ also depends on the cap $\bar{\alpha}$ through the insurance-market equilibrium; we suppress this argument except where the comparative static in $\bar{\alpha}$ is concerned).

Proposition 5 ((Equilibrium List Price)). *Suppose that*

1. $Q_k(w_A, w_B)$ *is twice continuously differentiable and the game is symmetric under relabeling the two drugs—*$Q_A(w_A, w_B) = Q_B(w_B, w_A)$ *for all* (w_A, w_B)*, with a common penalty* $\Lambda(\cdot)$ *(which implies* $Q_A(w, w) = Q_B(w, w)$ *on the diagonal);*

2. $\Phi(p, p) > 0$*;*

3. $\Phi(w_A, w_B)$ *is non-increasing in each argument, with the own effect dominating the cross effect (*$|\partial\Phi/\partial w_B| \geqslant |\partial\Phi/\partial w_A|$*);*

4. $\Lambda(\cdot)$ *is uniformly strictly convex (*$\Lambda'' \geqslant \underline{c} > 0$ *on* $[0, \bar{d}]$*) with* $\Lambda'(0) = 0$*, and* $\Lambda'(d) \to \infty$ *as* $d \to \bar{d}$ *for some upper bound* $\bar{d}$*, with* $\bar{w} \equiv p + \bar{d}$ *for the corresponding upper bound on the list price;*

Then the list-price-setting game has a unique Nash equilibrium at $w^\star > p$*, characterized by the first-order condition*

$$\Lambda'(w^\star - p) = \Phi(w^\star, w^\star) \tag{A2}$$

Moreover, if $\partial\Phi/\partial\bar{\alpha} > 0$*, then* $w^\star$ *is increasing in the coinsurance cap* $\bar{\alpha}$*.*

Proof. Manufacturer B's profit is

$$\pi_B(w_A, w_B) = p \cdot Q_B(w_A, w_B) - \Lambda(w_B - p).$$

1. *Existence and uniqueness.* The first-order condition for manufacturer B is

$$\frac{\partial \pi_B}{\partial w_B} = \Phi(w_A, w_B) - \Lambda'(w_B - p) = 0 \tag{A3}$$

where $\Phi(w_A, w_B) = p \cdot \partial Q_B/\partial w_B$. By condition (iii), $\Phi(w_A, w_B)$ is non-increasing in w_B for each w_A, and by condition (iv), $\Lambda'(\cdot)$ is strictly increasing. Hence the left-hand side of (A3) is strictly decreasing in w_B: manufacturer B's profit π_B is strictly concave in w_B for each w_A, giving at most one solution. Strict concavity

in own price rules out mixed-strategy equilibria: against any mixed strategy σ_A of the rival, the expected payoff $\int \pi_B(w_A, w_B)\, d\sigma_A(w_A)$ is a mixture of strictly concave functions and hence itself strictly concave in w_B, so manufacturer B has a unique pure best response to any σ_A and cannot randomize in equilibrium. Any symmetric fixed point must satisfy

$$\Lambda'(w^\star - p) = \Phi(w^\star, w^\star).$$

We first show this equation has a unique solution in $(p, \bar{w})$, then show the best-response map is a contraction, so this symmetric fixed point is the unique Nash equilibrium.

At a candidate symmetric equilibrium $w_A = w_B = w^\star$, the condition becomes $\Lambda'(w^\star - p) = \Phi(w^\star, w^\star)$.

At $w^\star = p$: the left-hand side is $\Lambda'(0) = 0$ while the right-hand side is $\Phi(p, p) > 0$ by condition (ii). At $w^\star \to \bar{w}$ (where $\Lambda'(\bar{w} - p) \to \infty$ by condition (iv)): the left-hand side exceeds the right-hand side (since Φ is bounded by smoothness of Q_B). By continuity of both sides and the intermediate value theorem, a solution $w^\star \in (p, \bar{w})$ exists. Since $\Lambda'(\cdot)$ is strictly increasing and $\Phi(w, w)$ is non-increasing in w along the diagonal (by condition (iii), both partial derivatives of Φ are non-positive), the symmetric solution is unique.

Finally, we rule out asymmetric pure-strategy equilibria. Implicit differentiation of (A3) gives the best-response slope

$$\frac{dw_B^\star}{dw_A} = \frac{\partial\Phi/\partial w_A}{\Lambda''(w_B - p) - \partial\Phi/\partial w_B},$$

whose absolute value is uniformly below one. Since the best response is continuous, piecewise continuously differentiable, has slope zero on any lower-corner region, and has derivative bounded in absolute value by $\rho < 1$ wherever interior, the mean-value theorem applied piecewise implies that each best response is globally Lipschitz with modulus ρ. By the uniform strict convexity of Λ (condition (iv)), $\Lambda'' \geqslant \underline{c} > 0$, so the denominator satisfies $\Lambda'' + |\partial\Phi/\partial w_B| \geqslant \underline{c} + |\partial\Phi/\partial w_A|$ using the own-dominance in condition (iii). Hence $|dw_B^\star/dw_A| \leqslant M/(\underline{c} + M) < 1$, where $M = \sup|\partial\Phi/\partial w_A| < \infty$ on the compact set $[p, \bar{w}]^2$ (by smoothness of Q_B).

This bound governs the interior of the best response. Strict concavity of π_B in w_B—from Φ non-increasing in w_B (condition (iii)) and the strict convexity of Λ—

makes the best response $w_B^\star(w_A)$ single-valued and continuous. Where it is interior it obeys the first-order condition and the slope bound above; at the lower corner $w_B = p$ it is locally constant (slope 0); and the upper boundary is never optimal, since $\Lambda'(w_B - p) \to \infty$ as $w_B \to \bar{w}$ forces $\partial\pi_B/\partial w_B < 0$. The joint best-response map is therefore a single-valued contraction on $[p, \bar{w}]^2$ with modulus $M/(\underline{c} + M) < 1$, and admits a unique fixed point; by symmetry of the game (condition (i)), that fixed point is the symmetric profile $(w^\star, w^\star)$.

2. *Comparative static.* Apply the implicit function theorem to the equilibrium condition $G(w^\star, \bar{\alpha}) \equiv \Lambda'(w^\star - p) - \Phi(w^\star, w^\star; \bar{\alpha}) = 0$:

$$\frac{dw^\star}{d\bar{\alpha}} = -\frac{\partial G/\partial\bar{\alpha}}{\partial G/\partial w^\star} = -\frac{-\partial\Phi/\partial\bar{\alpha}}{\Lambda''(w^\star - p) - \left.\frac{d\Phi(w,w)}{dw}\right|_{w=w^\star}}.$$

The denominator is positive: $\Lambda'' > 0$ by condition (iv), and $d\Phi(w,w)/dw = \partial\Phi/\partial w_A + \partial\Phi/\partial w_B \leqslant 0$ on the diagonal since condition (iii) makes Φ non-increasing in each argument. The numerator has sign $\partial\Phi/\partial\bar{\alpha} > 0$ under the additional assumption. Hence $dw^\star/d\bar{\alpha} > 0$.

The conditions in the proposition are stated as reduced-form properties of the expected-fills function $Q_k(\cdot)$ and the penalty $\Lambda(\cdot)$, without committing to a particular model of patient demand for drugs. Condition (ii) requires that a marginal list-price increase generates a positive revenue gain. Intuitively, when a manufacturer raises its list price slightly, plans covering its drug can offer lower premiums. Patients who weight premiums heavily sort toward these lower-premium plans, generating a positive demand response. Condition (iii) ensures diminishing returns from further list-price increases for each fixed rival price, which together with the strict convexity of $\Lambda(\cdot)$ guarantees strict concavity of each manufacturer's objective in its own list price. This yields single-valued best responses, ruling out mixed-strategy equilibria. Finally, the dominance of own over cross effects in condition (iii) makes the best-response map a contraction, which rules out asymmetric pure-strategy equilibria as well.

A leading example is Hotelling demand: patients are uniformly distributed on $[0, 1]$ with transport cost $t > 0$, drug A at 0 and drug B at 1, choosing among plans (and thus drugs) after observing premiums and effective cost-sharing. In this case, $Q_B(\cdot)$ varies smoothly in w_B through the premium channel described above, and a sufficient condition for (ii) is $\kappa\mu^\star > 1$, where $\mu^\star$ is the average risk

type in the pool; we omit the derivation.[39]

□

A.6 List Price Passthrough without Adverse Selection

The main text focuses on the adverse selection channel, but the passthrough prediction can also be consistent with a model without adverse selection and just patients overweighting premiums.

Suppose patient risk types θ are *observable* to plans, so each type receives a type-specific contract. All other primitives are as in Section 2. Plans choose coinsurance $\alpha \in [0, \bar{\alpha}]$ with $\bar{\alpha} < 1$, the drug has net price p and list price $w \geqslant p$, and free entry imposes zero profit on every contract. A plan serving type θ therefore charges

$$r(\theta) = \theta(p - \alpha w),$$

and the patient's utility on this contract is

$$U(\theta) = (1-\theta) \cdot u(W - \kappa r) + \theta \cdot u(W - \kappa r - \alpha w).$$

Substituting the zero-profit premium and writing $c = \alpha w$ for effective cost-sharing, the competitive contract maximizes U over $c \in [0, \bar{\alpha} w]$ subject to $r = \theta(p - c) \geqslant 0$.

Differentiating along the zero-profit locus with $a = u'(W - \kappa r)$ and $b = u'(W - \kappa r - c)$ ($b > a$ by strict concavity):

$$\frac{dU}{dc} = \theta[\kappa(1-\theta)a + (\kappa\theta - 1)b]. \tag{A4}$$

At $\kappa = 1$, this reduces to $\theta(1-\theta)(a-b) < 0$: patients prefer lower cost-sharing, as in the standard model. At $\kappa > 1$, the premium-overweighting channel raises both terms. In particular, $dU/dc > 0$ whenever

$$\kappa > \frac{b}{(1-\theta)a + \theta b} \equiv \kappa^{\star}(\theta), \tag{A5}$$

where $\kappa^{\star}(\theta) > 1$ since $b > a$. With $\kappa \in [2, 5]$ (Abaluck and Gruber, 2011), this threshold is exceeded for all but the lowest-risk types under empirically standard preferences, so the competitive contract sets $\alpha = \bar{\alpha}$ for the bulk of the type distribution.

[39] With $\kappa \in [2, 5]$ from Abaluck and Gruber (2011), this is satisfied for plausible parameter values.

When $\alpha = \bar{\alpha}$ binds and $w \in (p, p/\bar{\alpha})$, effective cost-sharing is $c = \bar{\alpha}w$ and the zero-profit premium is $r = \theta(p - \bar{\alpha}w) \geqslant 0$. Raising the list price by dw increases out-of-pocket costs by $\bar{\alpha} \cdot dw$, the same passthrough as in Proposition 3. In contrast, unconstrained plans (no cap on α) can adjust α downward to hold c fixed, so list prices remain irrelevant by Proposition 1. The channel shuts off at $w \geqslant p/\bar{\alpha}$, where $r = 0$ and further list-price increases cannot be passed through without violating the non-negative-premium constraint.

The mechanism differs from the main text: the cap binds not because adverse selection drives cream-skimming, but because κ-distorted patients prefer low-premium, high-cost-sharing contracts and competitive plans cater to this preference. The empirical prediction, however, remains identical. List-price increases pass through to patient out-of-pocket costs in cap-constrained plans and not in unconstrained plans.

B Additional Background

B.1 The Medicaid Drug Rebate Program

Here, we reproduce Figure 2 of Feng et al. (2023) to illustrate the incentives created by the combination of the basic rebate and the inflation penalty clauses in Medicaid. The scenario represented is the evolution of the Medicaid reimbursement price for a hypothetical drug with a launch list price of $100. In this scenario, the drug's price has doubled since launch, whereas CPI has only increased by 10 percent. As a result, the current Medicaid price is given by $200-$46.2-$90=$63.8, which is lower than the initial Medicaid price of $76.9. This highlights the strong incentives against increasing list price beyond the rate of inflation.

One issue to note is that the inflation penalty introduced as part of the Inflation Reduction Act (IRA) of 2022 creates much weaker incentives. Unlike Medicaid, there is no mandatory rebate rate in Medicare, so the inflation penalty only serves to neutralize any list price increase above the rate of inflation. Of course, Medicare Part B and Medicare Part D prices depend on rebate negotiations by private entities (for Part B through referencing and for Part D through direct negotiations by private insurers and pharmacy benefit managers). If negotiated rebates tend to be a fixed percentage of list price, then the provision would generate similar incentives to those created by the MDRP.

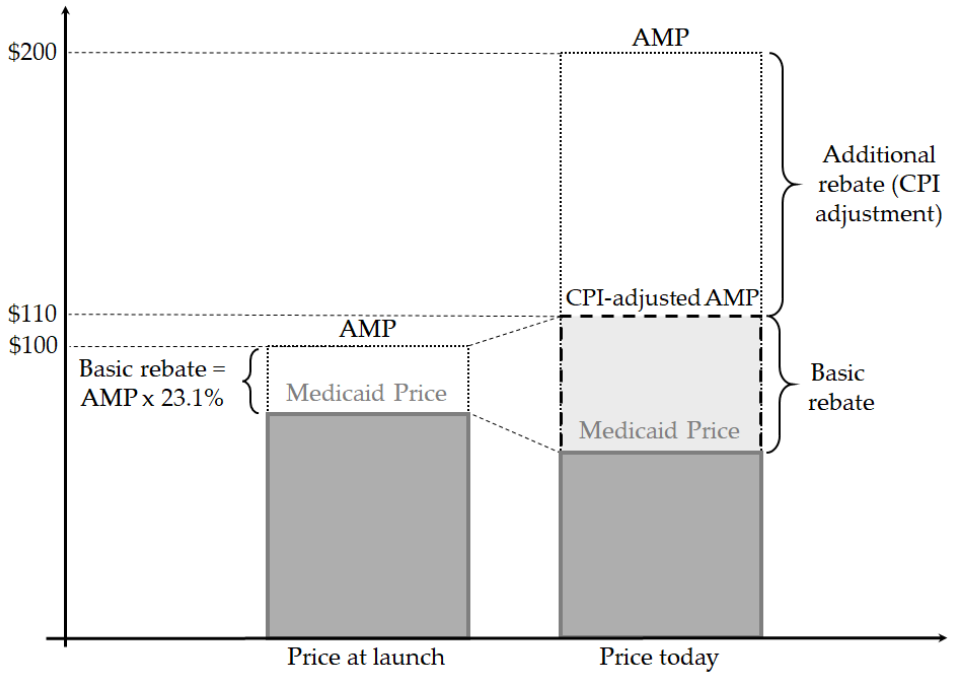

Figure A1: Impact of MDRP on changing list price

Notes: Reproduction of Figure 2 in Feng et al. (2023).

C Additional Results and Robustness Checks

C.1 Drug sample

Table A1: Statistics Related to Sample Selection Criteria

Criterion	Number of excluded drugs	Number of remaining drugs
		941
Part B sales greater than 20%	157	784
Injectables	135	649
Oncology	44	605
Invoice sales less than net sales flags	256	349
Contraceptives and smoking cessation drugs	9	340
Products launched before 1992	12	328
Products with only one year of data	24	304

The table tracks the number of unique drugs remaining after each successive selection criterion. The last two rows implement the market-age restriction at the drug level: we drop products launched before 1992 and products with only one year of data, for which a price change cannot be computed. This yields a core sample of 304 drugs. As described in the main text, the fifteen-year restriction is additionally enforced at the drug-year level: within these 304 drugs we retain only drug-years in which the drug is fifteen or fewer years past launch, so a drug crossing the threshold mid-panel contributes its earlier years but not its later ones. These within-drug exclusions are not reflected in the drug counts above, which is why the estimation sample contains fewer drug-years than the 304-drug core sample alone would imply.

C.2 Market shares of plan types in the MarketScan Data

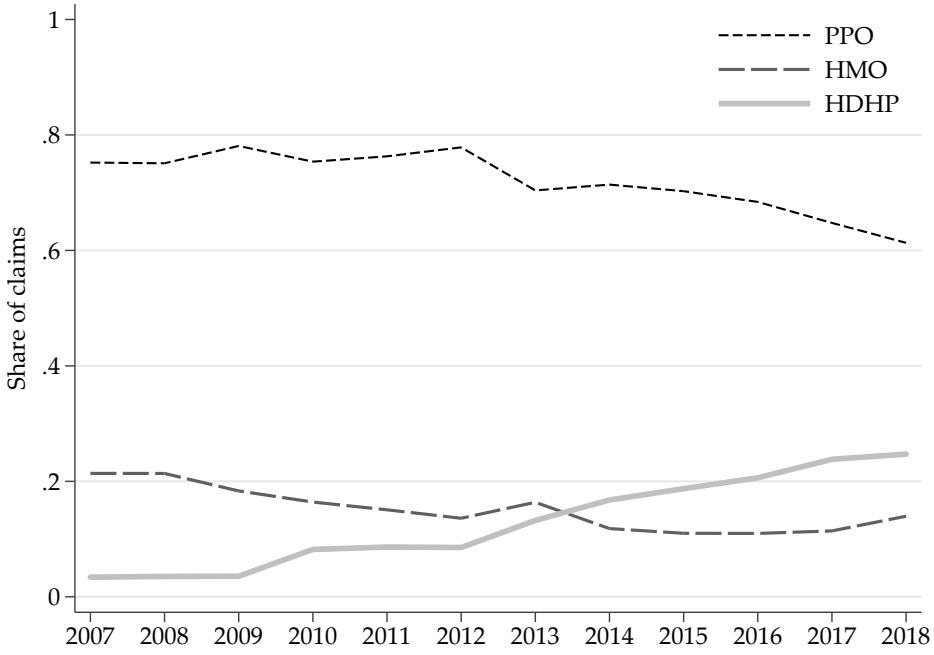

Figure A2: Market shares of plan types in the MarketScan Data

Notes: Plot of share of MarketScan claims from individuals in HMO, PPO, or HDHP plans from 2007 to 2018. A small number of other less common plan types are excluded from this calculation.

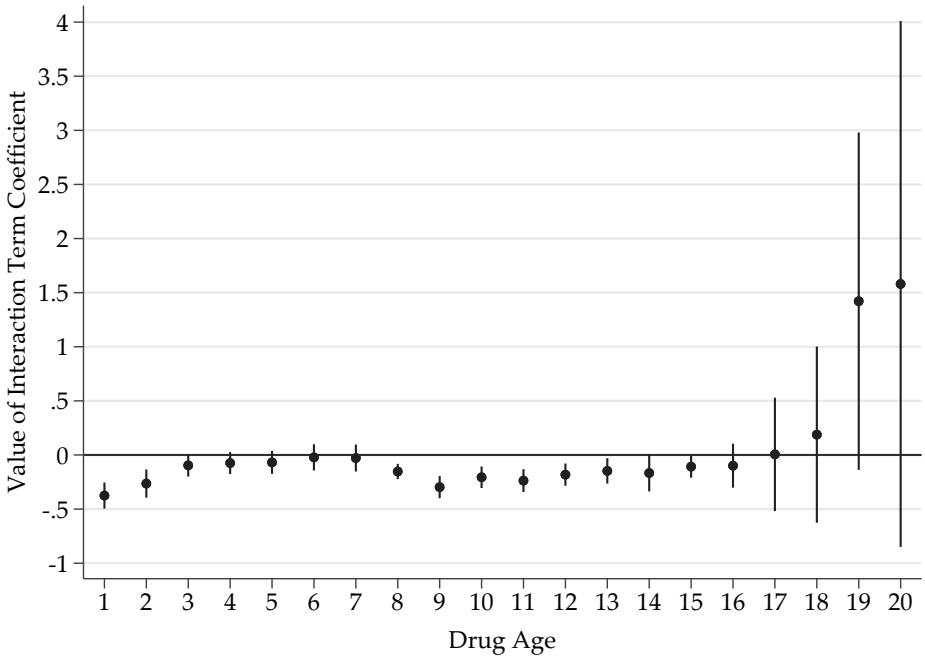

Figure A3: Plot of MMS×Drug-Age FE Coefficients from 1st Stage Model

Notes: Estimates for drug ages above 18 have very large standard errors due to small sample size. For the sake of scale, this plot is restricted to age 20 and under.

C.3 Estimates of First-Stage by Age of Drug

Figure A3 presents the estimates of Equation 4, broken down by the age of each drug. The estimates attempt to measure how responsive drugs are to the list price incentives created by the MDRP. As noted in the main text, drugs eventually hit the 100 percent rebate cap, after which the incentives to curb list price growth are nullified. Our estimates for older drugs are consistent with this story.

C.4 Results including change in net price and therapeutic area fixed effects as controls

We repeat our analysis but add change in net price and therapeutic area fixed effects as controls. We measure therapeutic area at the first level of the Anatomical Therapeutic Chemical classification system (ATC-1) developed by the World Health Organization.[40] There are 13 separate ATC-1s in our data. We find that the results are qualitatively sim-

[40]See, https://www.who.int/tools/atc-ddd-toolkit/atc-classification, accessed June 29th, 2026.

Table A2: Evidence Related to the Instrument (Net Price Control)

Panel A – First Stage Estimates

	List Price Growth	
	(1)	(2)
MMS	-0.142	-0.137
	(0.0324)	(0.0407)
$\Delta \ln(\text{net})$	0.119	0.117
	(0.0420)	(0.0424)
Year FE	X	X
ATC FE		X
N	1,582	1,574
R^2	0.173	0.194
Kleibergen-Paap F-stat	19.30	11.41

Notes: Estimates of Equation 4, using data from SSR Health on WAC to measure list price growth. Standard errors clustered at the drug level.

ilar, though noisier. This likely reflects a weaker first stage: because Medicaid Market Share is correlated within therapeutic area, the therapeutic area fixed effects absorb much of its variation, leaving less for the instrument to explain.

C.5 Comparison of NADAC and allowed amount increases

We study high-frequency changes in NADAC and allowed amount. NADAC is reported on a weekly basis. Allowed amount can also be calculated on a weekly basis.

Measurement issues: units are potentially not comparable across NADAC and allowed amount. Additionally, allowed amount can vary across pharmacies (e.g., if they are in network or out of network). To minimize these issues we restrict to pills or tablets (where we are reasonably sure that NADAC refers to a single pill and allowed amount per day supply is the comparable number), and to drugs with a sufficiently large number of fills in each week (100 or more). These restrictions limit this analysis to a handful of (potentially unrepresentative) NDCs.

We define a jump in NADAC as

- an increase in NADAC of 2% or more (we do not look at decreases because they are extremely rare)

Table A3: Response in out-of-pocket Costs – Commercial Plans (Net Price Control)

	$\Delta \log(OOP)$													
	All Plans								PPO		HMO		HDHP	
	All claims		Deductible Claims		Coinsurance Claims		Copay Claims		All claims		All claims		All claims	
Panel A. OLS Estimates														
$\Delta \ln(WAC)$	0.231	0.229	0.254	0.235	0.313	0.291	0.165	0.166	0.201	0.194	0.110	0.121	0.279	0.258
	(0.082)	(0.082)	(0.121)	(0.117)	(0.161)	(0.158)	(0.065)	(0.067)	(0.085)	(0.084)	(0.061)	(0.062)	(0.093)	(0.086)
$\Delta \ln(\text{net})$	-0.051	-0.051	-0.008	-0.014	0.058	0.058	-0.026	-0.024	-0.008	-0.005	-0.076	-0.067	-0.068	-0.071
	(0.031)	(0.031)	(0.027)	(0.027)	(0.019)	(0.019)	(0.024)	(0.024)	(0.019)	(0.019)	(0.040)	(0.039)	(0.025)	(0.025)
Panel B. Instrumental Variable Estimates														
$\Delta \ln(WAC)$	0.057	0.248	0.570	1.124	1.109	1.180	0.076	0.066	-0.025	-0.106	0.431	0.871	1.138	1.365
	(0.310)	(0.400)	(0.364)	(0.633)	(0.442)	(0.603)	(0.329)	(0.422)	(0.352)	(0.482)	(0.377)	(0.570)	(0.379)	(0.545)
$\Delta \ln(\text{net})$	-0.030	-0.053	-0.054	-0.138	-0.052	-0.063	-0.015	-0.013	0.020	0.031	-0.114	-0.155	-0.176	-0.207
	(0.042)	(0.058)	(0.062)	(0.107)	(0.088)	(0.108)	(0.041)	(0.054)	(0.045)	(0.063)	(0.058)	(0.089)	(0.077)	(0.099)
Kleibergen-Paap F-stat	19.61	11.58	18.65	11.62	19.11	12.39	19.61	11.58	19.48	11.27	18.19	10.65	20.87	13.99
Year FE	X	X	X	X	X	X	X	X	X	X	X	X	X	X
ATC FE		X		X		X		X		X		X		X
N	1,568	1,560	1,526	1,520	1,538	1,530	1,568	1,560	1,566	1,558	1,541	1,533	1,540	1,532

Notes: Estimates of Equation 3 using claims data from MarketScan. Standard errors clustered at the drug level.

Table A4: Response in out-of-pocket Costs – Medicare Plans (Net Price Control)

	$\Delta \log(OOP)$					
	All Plans		Basic Plans		Enhanced Plans	
PANEL A. OLS ESTIMATES						
$\Delta \ln(WAC)$	0.245	0.228	0.208	0.184	0.209	0.192
	(0.119)	(0.117)	(0.100)	(0.093)	(0.122)	(0.120)
$\Delta \ln(\text{net})$	0.013	0.013	0.037	0.040	-0.004	-0.004
	(0.016)	(0.016)	(0.021)	(0.022)	(0.018)	(0.019)
PANEL B. INSTRUMENTAL VARIABLE ESTIMATES						
$\Delta \ln(WAC)$	0.942	1.229	1.075	0.723	0.558	0.769
	(0.238)	(0.435)	(0.466)	(1.017)	(0.219)	(0.415)
$\Delta \ln(\text{net})$	-0.124	-0.179	-0.161	-0.081	-0.084	-0.133
	(0.099)	(0.148)	(0.164)	(0.254)	(0.085)	(0.134)
Kleibergen-Paap F-stat	16.06	8.10	9.17	2.97	14.02	6.10
ATC FE		X		X		X
Year FE	X	X	X	X	X	X
N	1,410	1,408	1,350	1,348	1,383	1,381

Notes: Estimates of Equation 3 using CMS's 20% random sample of Medicare Part D beneficiaries. Basic plans follow the standard Medicare benefit design or offer actuarially equivalent benefits. Enhanced plans are more generous, and offer additional coverage. Standard errors clustered at the drug level. The first stage is weak in the Medicare sample once therapeutic-area (ATC) fixed effects are included (Kleibergen-Paap F between 3 and 8), so the ATC-FE columns should be interpreted with caution.

Table A5: Changes in Allowed Amount (Net Price Control)

	Commercial		Medicare	
PANEL A. OLS ESTIMATES				
$\Delta \ln(WAC)$	0.571	0.558	0.486	0.464
	(0.212)	(0.216)	(0.206)	(0.205)
$\Delta \ln(\text{net})$	-0.015	-0.013	0.006	0.007
	(0.018)	(0.018)	(0.017)	(0.017)
PANEL B. INSTRUMENTAL VARIABLE ESTIMATES				
$\Delta \ln(WAC)$	1.180	1.237	1.186	1.277
	(0.251)	(0.332)	(0.255)	(0.459)
$\Delta \ln(\text{net})$	-0.087	-0.092	-0.131	-0.149
	(0.062)	(0.071)	(0.116)	(0.152)
Kleibergen-Paap F-stat	19.61	11.58	16.06	8.10
ATC FE		X		X
Year FE	X	X	X	X
N	1,568	1,560	1,410	1,408

Notes: Estimates of Equation 3, using allowed amount as the outcome. Commercial estimates are based on MarketScan claims and Medicare estimates are based on CMS's 20% random sample of Medicare Part D beneficiaries. Standard errors clustered at the drug level. In the Medicare sample, the first stage is weak once therapeutic-area (ATC) fixed effects are included (Kleibergen-Paap F = 8.1), so that column should be interpreted with caution.

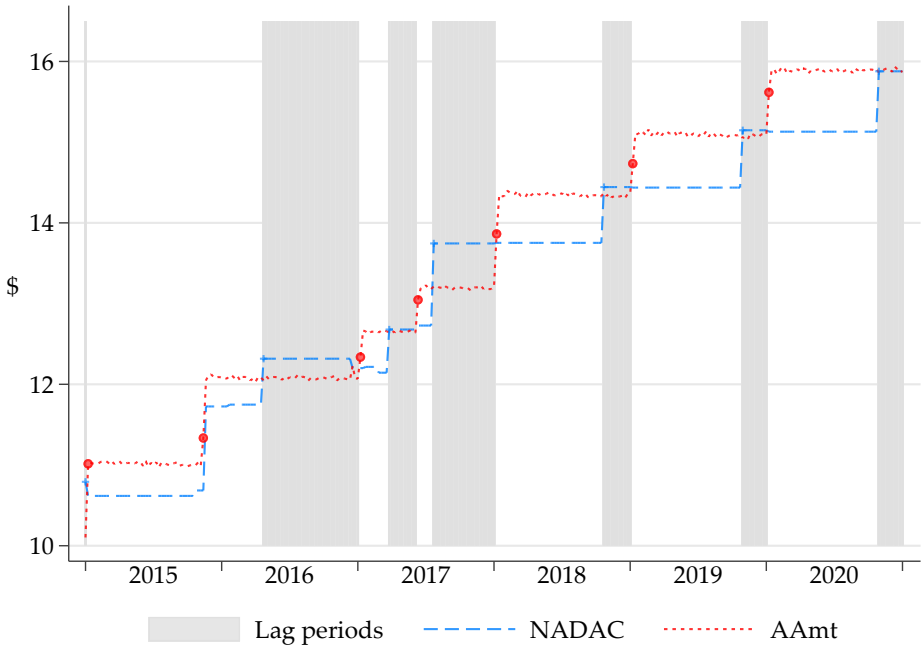

Figure A4: NADAC vs. Allowed Amount in Medicare claims for Januvia 50 MG tablet (NDC: 00006011231)

Notes: Plot of NADAC and average per-claim Medicare allowed amount at the weekly level for NDC 00006011231. Periods shaded in grey represent gaps between changes in NADAC and changes in allowed amount.

- that is not preceded by an increase in allowed amount in the previous 4 weeks (we do this in case NADAC increases are reported with a lag)
- and that is permanent (i.e. does not return back to baseline within 4 weeks)

We then identify increases in allowed amount following these jumps in NADAC by applying a similar algorithm *starting with the week in which NADAC jumped*. The time between a NADAC jump and an allowed amount jump is considered a "lag." Figure A4 shows an example of the patterns identified by this algorithm for one of the drugs in the sample.

Figure A5 shows the distribution of lags after each increase in NADAC. We find that in around 60% of cases, the jump is immediately (i.e., in the same week) followed by a jump in allowed amount. However, in the remaining 40% of cases, we find a lag. The average lag (conditional on existing) is about 12 weeks.

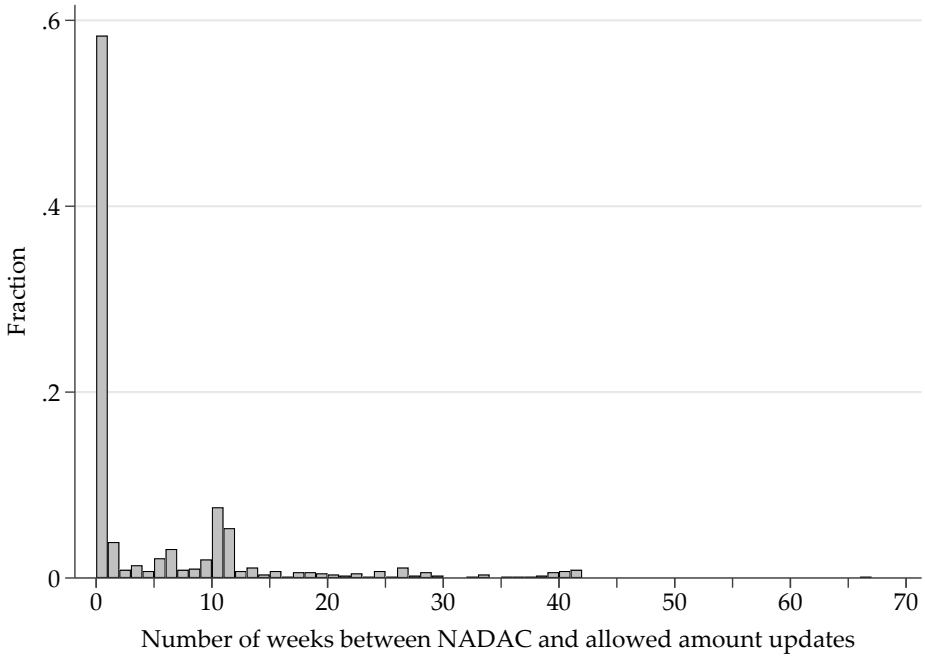

Figure A5: Lag between NADAC and Allowed Amount updates

Notes: Plot of number of weeks between successive NADAC and allowed amount increases.