The Price No One Shows You

Contracting over Pharmaceutical Formularies and Rebates — Kate Ho & Robin S. Lee (*Econometrica*, 2026)

Felipe Ovando Turra

Roadmap

The empirical problem of drug prices

A huge market — with an invisible price

- U.S. retail drug spending: **$467 billion** in 2024 — **9%** of total health expenditures, growing ∼**8% per year** since 2020.
- Branded manufacturers post **high list prices** under patent protection.
- But nobody pays the list price. Manufacturers pay back **rebates** to intermediaries — post-sale payments that cut the *net* price.
- Rebates are **confidential**. They are the price that actually clears the market, and *no researcher observes them*.

The methodological challenge

How do you study a market when the **equilibrium price is unobservable**?

Four questions

1. Can you **measure a price** that nobody lets you see?
2. Why would a manufacturer pay for where its **rival** is placed?
3. Where does the bargaining power of a middleman who only says *"take it or leave it"* come from?
4. What does this tell us about what **you** pay at the pharmacy?

The answer to (1) is the method: impose the equilibrium structure of a contracting game on what we *do* observe.

Four questions

1. Can you **measure a price** that nobody lets you see?
2. Why would a manufacturer pay for where its **rival** is placed?
3. Where does the bargaining power of a middleman who only says *"take it or leave it"* come from?
4. What does this tell us about what **you** pay at the pharmacy?

The answer to (1) is the method: impose the equilibrium structure of a contracting game on what we *do* observe.

Four questions

1. Can you **measure a price** that nobody lets you see?
2. Why would a manufacturer pay for where its **rival** is placed?
3. Where does the bargaining power of a middleman who only says *"take it or leave it"* come from?
4. What does this tell us about what **you** pay at the pharmacy?

The answer to (1) is the method: impose the equilibrium structure of a contracting game on what we *do* observe.

Four questions

1. Can you **measure a price** that nobody lets you see?
2. Why would a manufacturer pay for where its **rival** is placed?
3. Where does the bargaining power of a middleman who only says *"take it or leave it"* come from?
4. What does this tell us about what **you** pay at the pharmacy?

The answer to (1) is the method: impose the equilibrium structure of a contracting game on what we *do* observe.

Four questions

1. Can you **measure a price** that nobody lets you see?
2. Why would a manufacturer pay for where its **rival** is placed?
3. Where does the bargaining power of a middleman who only says *"take it or leave it"* come from?
4. What does this tell us about what **you** pay at the pharmacy?

The answer to (1) is the method: impose the equilibrium structure of a contracting game on what we *do* observe.

Four questions

1. Can you **measure a price** that nobody lets you see?
2. Why would a manufacturer pay for where its **rival** is placed?
3. Where does the bargaining power of a middleman who only says *"take it or leave it"* come from?
4. What does this tell us about what **you** pay at the pharmacy?

The answer to (1) is the method: impose the equilibrium structure of a contracting game on what we *do* observe.

Why this paper? An identification strategy, not just a drug study

Two contributions

1. **Theory.** A model of *multidimensional* contracting: rebates condition on the **entire formulary** — a drug's own tier *and its rivals'*. Nearly all prior vertical-contracting work lets contracts condition only on a product's *own* inclusion (Crawford–Yurukoglu, Grennan, Gowrisankaran–Nevo–Town, Ho–Lee 2017/2019).
2. **Method.** Map **estimated demand responses** + **observed list prices** $\Rightarrow$ **implied equilibrium rebates**. *No rebate data required.*

The transferable idea

Nothing here is specific to drugs: any market with **menu contracts** and **cross-firm externalities** is fair game.

Who is who in this market

The players that *decide*

- **Manufacturers.** Patent-protected; set **list prices** unilaterally, at supracompetitive levels.
- **Plan sponsor.** Bears the financial risk of its enrollees' drugs. Here: **Princeton University** (self-insured employer).
- **PBM.** The sponsor lacks the scale and expertise, so it **delegates**: the PBM (i) designs the **formulary**, (ii) **negotiates rebates**, (iii) administers claims.

The players that are *plumbing*

- **Wholesalers & pharmacies.** Move the pills, collect a **fixed margin** (dispensing fees). Not strategic — abstracted away.
- **Consumers.** Choose drugs given their **copay**. Their demand is the terrain the game is fought on.

The art of the model

Everything that does not *choose* becomes a constant.

The formulary: the PBM's instrument

A formulary is an *assignment*: each drug → a tier → a copay

Tier	Your copay (30-day retail)	Effect on demand
Generic	\$5	cheap ⇒ most people switch here
Preferred (P)	**\$25**	moderate
Non-preferred (NP)	**\$40**	costly ⇒ demand flows away
Excluded (E) / MPTD	full list price	prohibitive

The causal chain — *this is the thesis of the whole paper*

tier ⇒ **copay** ⇒ **volume** ⇒ a manufacturer will **pay** not to be demoted.

That payment is the rebate. The formulary is the **shelf plan**; the rebate is what Coca-Cola pays for eye level.

Following the money: two moments, one fill of Crestor

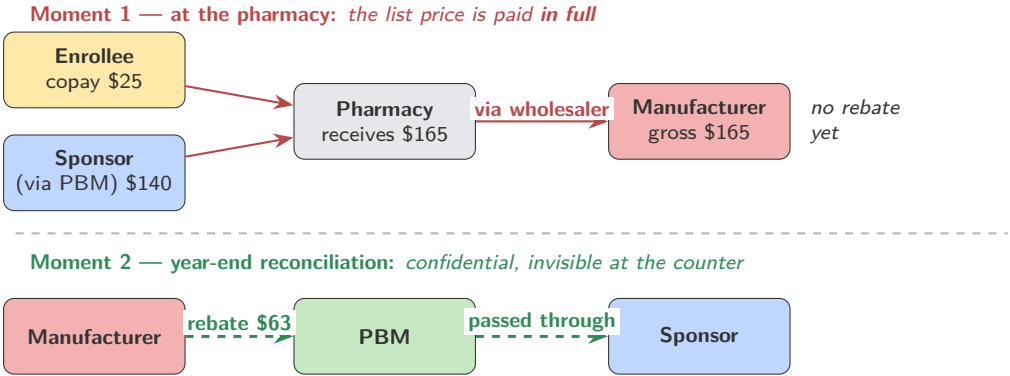

*The paper observes copay + PBM payment (⇒ list price). It does **not** observe the rebate — and reconstructs it.*

What is observable, and what is not

Object	Magnitude	Who sees it	In the data?
List price	∼\$165	public	**Yes** (copay + PBM payment)
Copay	\$25 (tier-set)	enrollee	**Yes**
Sponsor payment	∼\$140	sponsor	**Yes**
Rebate	∼\$63 (38%)	confidential	**No**
Net price	∼\$102	confidential	**No**

The frontier the paper crosses

Everything at the counter is visible. Everything after it is dark.
The contribution is to **illuminate the dark half using the lit one** + a demand model.

Setup and timing

Primitives

One PBM, M branded manufacturers.
A formulary $\mathbb{F} = \{\mathbb{F}_1, \ldots, \mathbb{F}_M\}$ with $\mathbb{F}_i \in \{P, NP, E\}$; the set of formularies is $\mathcal{F} \equiv \{P, NP, E\}^M$.

$$\mathbb{C}_i : \mathcal{F} \to \mathbb{R}_0^+$$

(rebate contract of manufacturer i)

Three stages

1. PBM makes **private, take-it-or-leave-it** offers $\mathbb{C}_i$; manufacturers accept/reject.
2. Given accepted contracts, PBM **chooses the formulary** $\mathbb{F}$.
3. Given $\mathbb{F}$, **consumers choose** drugs.

No commitment — and this is what gives the model teeth

The PBM **cannot commit** in Stage 1 to a Stage-2 formulary: whatever it picks must be *credible* (optimal) given the contracts actually signed.

Solution concept: WPBE with **passive beliefs** (Hart–Tirole 1990; McAfee–Schwartz 1994).

The rebate contract is a *menu*, not a number

The apparent paradox

Contracts are signed in Stage 1, but $\mathbb{C}_i(\cdot)$ is a function *of the formulary* — chosen in Stage 2. How can the rebate depend on something that hasn't happened?

Resolution: it is an insurance policy, not a payment

- Stage 1 fixes the **entire contingent table**: "if the formulary is $\mathbb{F}'$, I pay $\mathbb{C}_i(\mathbb{F}')$; if it is $\mathbb{F}''$, I pay $\mathbb{C}_i(\mathbb{F}'')$; ..."
- Stage 2 merely **selects which entry is activated**.
- A manufacturer *wants* this: it will pay a lot to sit in P with its rival in NP, and nothing at all if it is demoted. Only a menu can express that.

A puzzle to keep in mind

Most entries of the table are for formularies that will **never be chosen**.
So why would anyone bother negotiating them? **Hold that thought.**

Payoffs

The PBM (acting for the sponsor)

$$\Pi_P(\mathbb{F}, \mathbb{C}) = \underbrace{\rho\, W(\mathbb{F})}_{\text{consumer welfare}} - \sum_i D_i(\mathbb{F}) \times \Big(\underbrace{p_i}_{\text{list}} - \underbrace{p^{oop}(\mathbb{F}_i)}_{\text{copay}} - \underbrace{\mathbb{C}_i(\mathbb{F})}_{\text{rebate}} \Big)$$

$\rho \geq 0$ weights enrollee welfare: it encodes **what type of sponsor** this is. Simulated at $\rho \in \{0, 1\}$, *not estimated*.

Manufacturer i (zero marginal cost)

$$\Pi_i\Big(\mathbb{F}, \mathbb{C}_i(\mathbb{F})\Big) = \underbrace{D_i(\mathbb{F})}_{\text{volume}} \times \Big(\underbrace{p_i - \mathbb{C}_i(\mathbb{F})}_{\text{net price}} \Big)$$

The copay does not appear — it already did its job *inside* $D_i(\mathbb{F})$.

What the empirics must deliver

p_i observed. $\mathbb{C}_i$ unknown. **The only thing we need from the data is $D_i(\mathbb{F})$.**

A toy game (two manufacturers, two tiers)

Payoffs *before* rebates

Formulary (M_1, M_2)	π_1	π_2
(P, P)	10	10
(P, NP)	16	6
(NP, P)	6	12
(NP, NP)	8	8

The PBM earns **zero** before rebates: it only collects what it extracts.

Reading the table

- Best for a firm: **alone in P**, rival in NP $\Rightarrow$ captures volume.
- Worst: **demoted** while the rival is promoted.
- The **generic** is the silent outside option that gives the demotion its bite.

The atomic logic — and the notation

$$\text{rebate } \mathbb{C}_i(\hat{\mathbb{F}}) = \underbrace{\pi_i(\hat{\mathbb{F}})}_{\text{profit in equilibrium}} - \underbrace{\Pi_i^d}_{\text{profit under your } \textbf{punishment}}$$

The **punishment** (formally, the *disagreement payoff*) is what the PBM would pick *if you rejected*.

Step 1: which formulary is chosen?

The equilibrium formulary maximizes *joint* surplus (Prop. 3 — *why*, in two slides)

Formulary	(P, P)	$(\mathbf{P}, \mathbf{NP})$	(NP, P)	(NP, NP)
$\pi_1 + \pi_2$	20	**22**	18	16

Intuition

The PBM ends up appropriating (via rebates) what the manufacturers are willing to pay.
So it wants the formulary that makes the **pie biggest** — more to extract.

Equilibrium formulary: $\mathbb{F}^* = (P, NP)$

M_1 in Preferred (earns 16 gross), M_2 in Non-preferred (earns 6 gross).
Now: how much does each one pay?

Step 2: two candidate equilibria — *same* formulary, *different* rebates

Equilibrium A

Contracts (total payments, ordered over $(P,P),(P,NP),(NP,P),(NP,NP)$):

$$\mathbb{C}_1 = (2, \mathbf{8}, 0, 0), \qquad \mathbb{C}_2 = (0, 0, \mathbf{0}, 2)$$

- If M_1 rejects: PBM's best collection from M_2 is **2**, at (NP, NP).
- $\Rightarrow$ **punishment** $= (NP, NP)$, where M_1 earns **8**.
- Rebate: $16 - 8 = \mathbf{8}$. PBM collects **8**.

Equilibrium B

$$\mathbb{C}_1 = (4, \mathbf{10}, 0, 2), \qquad \mathbb{C}_2 = (4, 0, \mathbf{6}, 2)$$

- If M_1 rejects: PBM's best collection from M_2 is **6**, at (NP, P).
- $\Rightarrow$ **punishment** $= (NP, P)$, where M_1 earns only **6**.
- Rebate: $16 - 6 = \mathbf{10}$. PBM collects **10**.

Remember the puzzle? Here is the answer.

Both satisfy the necessary conditions. Both are equilibria. **This is the multiplicity.**
The *only* difference is $\mathbb{C}_2$ at (NP, P) — **a formulary that is never chosen**.
Off-equilibrium entries are what make the threat credible — and hence what M_1 pays.

Step 3: the MI refinement selects one

Manufacturer Indifference (Definition 1)

For every $\mathbb{F}' \neq \hat{\mathbb{F}}$, set $\mathbb{C}_i(\mathbb{F}') = \max\{X, 0\}$ where X solves $\Pi_i(\mathbb{F}', X) = \Pi_i(\hat{\mathbb{F}}, \hat{\mathbb{C}}_i(\hat{\mathbb{F}}))$:

each manufacturer is held to the ***same payoff under every formulary****.*

Check Equilibrium B: both firms net **6** everywhere

	(P, P)	(P, NP)	(NP, P)	(NP, NP)
M_1: $\pi_1 - \mathbb{C}_1$	$10 - 4 = 6$	$16 - 10 = 6$	$6 - 0 = 6$	$8 - 2 = 6$
M_2: $\pi_2 - \mathbb{C}_2$	$10 - 4 = 6$	$6 - 0 = 6$	$12 - 6 = 6$	$8 - 2 = 6$

Equilibrium A **fails**: M_1 nets $8, 8, 6, 8$ — *not* indifferent.

Why MI $\Rightarrow$ efficiency (the promised "why" from Step 1)

If firms are **indifferent** across formularies, *all* the variation in joint surplus falls on the PBM.
It therefore **internalizes** it: its selfish choice *is* the efficient one. **That is Prop. 3.**

The economics of MI, and what it buys us

The economic content of the refinement

MI is the equilibrium in which the PBM **plays the manufacturers off against each other** — even in a *one-shot* game.

Two further results

- **Bernheim–Whinston.** If manufacturers *bid* instead (a **menu auction**), MI coincides exactly with **Truthful Nash Equilibria** — Coalition-Proof, hence "focal."
  Empirically the two games are close (28–38% vs. 25–33%): disagreement payoffs are substantial either way.
- **Computable.** An algorithm searches over profiles of perceived punishments; in the empirical application the MI-equilibrium is always **unique**.

Deferred to the appendix

A **no-multilateral-deviation** condition rules out the PBM "tricking" several manufacturers at once by exploiting passive beliefs. It closes sufficiency.

Data: Princeton University, 2011–2018

A clean laboratory

- A **single self-insured sponsor** with a **single PBM**: Express Scripts (2011–15), Optum (2016–18). No sponsor heterogeneity by construction.
- Focus: **statins**. Three major branded drugs (Lipitor, Crestor, Vytorin) + generics. **1,965** enrollees; **35,220** fills.
- Enrollees pay **fixed-dollar copays** — so the list price is *invisible* to their choice. (This is what licenses omitting price from demand.)
- List price per fill = patient payment + PBM payment. **Observed.**

The identification challenge

All three brands **move across tiers** — but mostly *contemporaneously with generic entry*.
How do we separate the **tier effect** from the **generic-entry effect**?

The clean variation: Vytorin, 2016

One drug is demoted *without* a generic

- In **2016** Princeton switches to Optum; **Vytorin is moved to Tier 3** ($40 copay).
- Its generic (ezetimibe/simvastatin) does **not** arrive until **2017**.
- Share falls from ∼**1.7%** (2015) to ∼**0.5%** (2016).

A tier move decoupled from generic entry.

The other two, for contrast

- **Lipitor**: demoted in 2012 — *same year* its generic enters.
- **Crestor**: stays on the branded tier *until* its generic arrives (2016).

Two further identification assumptions

- In 2011–15 tier placement is set **mechanically** by branded/generic status — not by demand shocks.
- Generic-entry **timing** is plausibly exogenous (well-documented FDA approval delays).
  Robustness: dropping the 3 months around generic entry (advertising) leaves estimates unchanged.

Demand: a logit with inertia

Utility of enrollee k, drug i, channel c, fill f

$$u_{kic} = \underbrace{\gamma_{m(i)}}_{\text{molecule FE}} + \underbrace{\gamma_{\text{tier}(i,t),c}}_{\textbf{tier}\times\text{channel FE}} + \underbrace{\alpha_k \mathbf{1}(c = M)}_{\text{channel prefs}} + \underbrace{\beta^B_k b(i)}_{\text{brand pref (RC)}} + \beta^M \mathbf{1}(f{=}1, c{=}M) - \underbrace{SC(\cdot)}_{\text{switching costs}} + \varepsilon_{kic}$$

Estimated by **simulated maximum likelihood** over each enrollee's observed sequence of fills.

Two modelling choices worth defending

- **No price term.** Copays are fixed dollars, absorbed by the **tier×channel** fixed effects. We only need *share responses to tier moves* — not price elasticities.
- **But then how do we get dollars?** $W(\mathbb{F})$ must be in \$ to enter the PBM's objective. They **normalize using the \$15 copay gap** between Tier 2 and Tier 3. *Omitting price has downstream consequences.*

Estimates (\$) — and why random coefficients matter

Cost of switching molecule: with RC **\$32.4** $\longrightarrow$ without RC **\$46.8**

Tier effects are stable ($-$\$7.1 / $-$\$22.1). Strip out preference heterogeneity and it **reloads onto inertia**: people look "stuck" when they merely *prefer*.

Tier placement moves volume — own effects

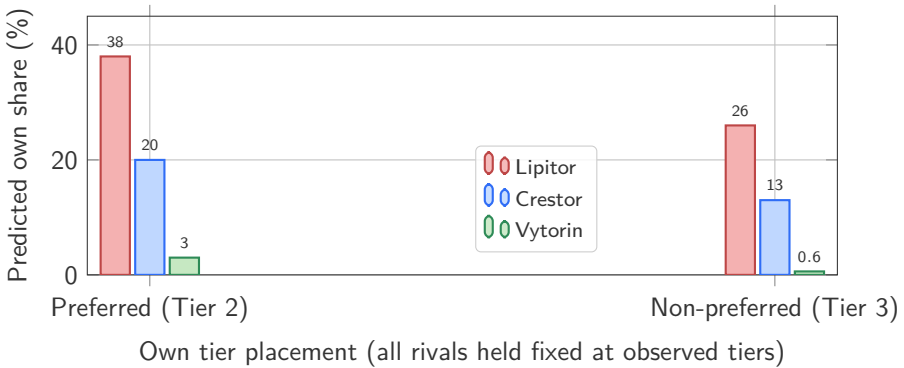

Why the predicted drops are *smaller* than the raw data

In the raw data Lipitor collapses from 37% to ∼2%. The model predicts only 38% → 26% — because it **isolates the tier effect**, holding generic entry fixed. *The raw collapse is tier + generic; the model separates them.*

. . . and cross effects: where does the volume *go*?

This is what makes the *rival-contingent* contract bite

The whole theory rests on rebates conditioning on the **rival's** tier. That only has force if demoting Lipitor **redirects volume toward Crestor** — not just toward generics.

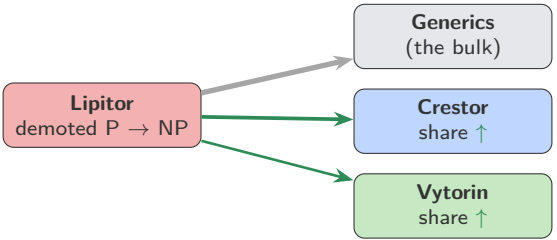

Fig. 1, right panel: as Lipitor's tier worsens, **rival branded shares rise**.

Cross-substitution exists ⇒ the rival's placement is worth paying for.

The asymmetry this creates

The **generic** absorbs most of the diverted volume ⇒ demand effects are **asymmetric**: being demoted P→NP (rival fixed at P) hurts *more* than staying at NP while the rival is promoted. **Remember this: it is why Crestor will pay 38% and Lipitor only 28%.**

The benchmark that makes everything interpretable

With a *single* branded tier and no exclusion . . .

The model predicts *zero* rebates.

Why

A manufacturer will not pay anything if its drug is offered at the **same copay regardless**. No differentiation $\Rightarrow$ no threat $\Rightarrow$ no leverage $\Rightarrow$ no rebate.

The corollary that reframes the debate

Rebates are not a *discount* the PBM negotiates.
They are the **price of a credible threat**.

Table IV: equilibrium rebates (Panel B, two drugs)

	ρ	**Equilibrium** (Lipitor, Crestor)	**Punishment** formularies	**Rebates (% of list), with 95% CI** (Lipitor, Crestor)
No exclusion	0	(NP, NP)	(NP, P) / (P, NP)	(10.0%, 17.4%)
	1	$(\mathbf{P}, \mathbf{P})$	(NP, P) / (P, NP)	**(28.1%, 38.2%)** [22.8, 34.4], [32.2, 45.9]
With exclusion	0	(NP, NP)	(E, P) / (P, E)	(58.1%, 69.9%)
	1	(P, P)	(E, P) / (P, E)	(66.5%, 77.5%) [46.4, 81.6], [59.8, 89.0]

Reading the table

- Rebates are **always positive** and **always below 100%**: manufacturers keep some gains from trade *even facing take-it-or-leave-it offers*.
- $\rho = 0$: both brands demoted to steer to generics; the punishment is *barely worse* than equilibrium $\Rightarrow$ **weak leverage, small rebates**.
- $\rho = 1$: both brands in P; demotion is now *painful* $\Rightarrow$ **strong leverage, large rebates**.

The validation: predicted rebates vs. independent data

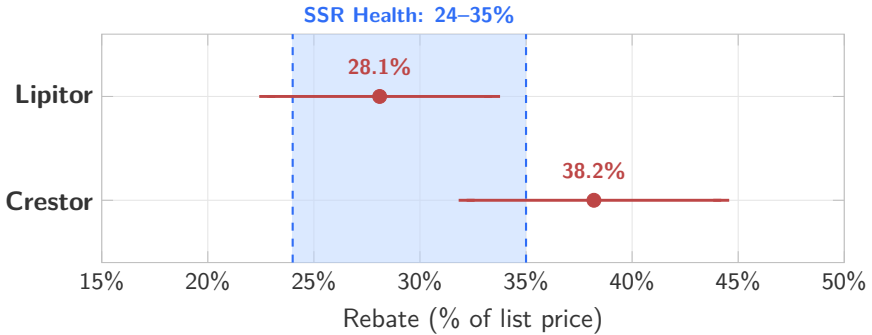

The headline result — with honest uncertainty

Red: model estimates + bootstrapped 95% CIs — built from **demand and list prices only**.

Blue band: rebates reported from **actual rebate data** (SSR Health).

Two independent routes to the same number. The intervals overlap.

Two further findings

1. Exclusion: huge rebates — but *calibrated*, not estimated

With exclusion as the threat, rebates jump to **58–78%**. **But exclusion is never observed in the data**, so its volume effect **cannot be estimated**. It is **assumed**: a \$55 copay, giving a **∼75% volume drop** (imported from Agha–Kim–Li). At those levels, Medicaid **MFN** ("best price") would also bind — and is abstracted away.

Two independent reasons to distrust that row of Table IV.

2. More flexibility can *hurt* the PBM

Adding an **intermediate tier** (\$32.50) sometimes *lowers* rebates: the PBM can no longer commit to a punishment harsh enough — it would rather use the intermediate tier upon disagreement, which weakens its leverage.

The threat is worth more than its execution.

What the model deliberately leaves out

Five abstractions, in order of bite

1. **List prices are fixed.** The most consequential. Manufacturers respond to higher rebates by *raising list prices* (documented for insulin).
2. ρ **is not estimated.** Formulary placement in the data is a mechanical response to generic entry, so PBM preferences are unidentified.
3. **Exclusion is calibrated**, not estimated — never observed in the data.
4. **PBM rents are assumed away.** Sponsor net payment = PBM net payment. The agency problem is not modelled.
5. **One sponsor, one PBM.** No PBM competition, no sponsor heterogeneity.

An honest scope

The paper answers "*given* list prices, what are the rebates?" — not "why are list prices so high?"

Take-home message

The unifying logic

A **credible threat** to move demand $\Rightarrow$ willingness to pay $\Rightarrow$ rebates.

With one branded tier: **zero rebates**. With P/NP: **28–38%**.
The PBM never had to *execute* the punishment — the threat did the work.

The method, not the statins, is the point

Structural equilibrium restrictions + observable demand = recovery of a confidential price.

- Confidential prices leave **behavioural fingerprints** in the demand system — if you know the game being played.
- The **template** generalizes: shelf space, insurer networks, carriage deals, slotting allowances.

Wherever a buyer can credibly redirect demand, look for the hidden rebate.

Thank you!

Questions?

Felipe Ovando Turra

Ho & Lee — "Contracting over Pharmaceutical Formularies and Rebates" (Econometrica, 2026)

Backup — Verifying the necessary conditions

Condition 1 (Lemma 1) — each manufacturer earns exactly its punishment payoff Π_i^d

	Equilibrium A		**Equilibrium B**	
	M_1	M_2	M_1	M_2
Punishment formulary	(NP, NP)	(P, NP)	(NP, P)	(P, NP)
Punishment payoff Π_i^d	8	6	6	6
Equilibrium net payoff	$16 - 8 = 8$ ✓	$6 - 0 = 6$ ✓	$16 - 10 = 6$ ✓	$6 - 0 = 6$ ✓

Condition 2 (eq. 6): $\Pi_P(\hat{\mathbb{F}}, \hat{\mathbb{C}}) \geq \Pi_P(\mathbb{F}', \mathbb{C}_{-i}^d) + \max\{\pi_i(\mathbb{F}') - \Pi_i^d, 0\}$

	Equilibrium A (LHS = 8)			**Equilibrium B** (LHS = 10)		
Deviate to:	(P, P)	(NP, P)	(NP, NP)	(P, P)	(NP, P)	(NP, NP)
RHS with M_1	2	0	2	8	6	4
RHS with M_2	6	6	2	8	6	4

All RHS $\leq$ LHS in **both** equilibria. ✓

BACKUP — Multilateral deviations (Appendix A.2)

Why the two conditions are not sufficient

They rule out deviations with **one** manufacturer at a time (a "delegated-agent" equilibrium). But the PBM can deviate with **several at once**, exploiting **passive beliefs**: each firm, receiving an odd offer, still believes its *rival* is on the equilibrium contract.

The "trick" (Equilibrium A of the toy game)

The PBM offers $\mathbb{C}_1' = (7.9, 8, 0, 0)$ and $\mathbb{C}_2' = (7.9, 0, 0, 0)$. Each firm accepts, believing (P, NP) will be chosen. But once *both* accept, the PBM prefers (P, P) and collects $7.9 + 7.9 = 15.8 > 8$.

What restores the equilibrium

A restriction on the contract space (eq. 17): no contract may leave a manufacturer, under *any* formulary, below its worst no-contract payoff $\underline{\Pi}_i \equiv \min_{\mathbb{F}} \Pi_i(\mathbb{F}, 0)$.
This caps the trick at $\mathbb{C}_1' = (2, \ldots)$, $\mathbb{C}_2' = (4, \ldots)$ — collecting only $6 < 8$. **No profitable multilateral deviation.**

Backup — Offer game vs. bidding game

	Equilibrium formulary	Rebates (Lipitor, Crestor)
Offer game (PBM makes offers)	(P, P)	(28.1%, 38.2%)
Bidding game (firms bid)	(P, P)	(24.7%, 32.9%)

Interpretation

- **Same formulary** in both — as Prop. 3 requires.
- Rebates are *higher* when the PBM makes offers: under bidding, each manufacturer captures part of the surplus its own inclusion generates.
- **But the gap is small**, which tells us that **disagreement payoffs are substantial** — manufacturers retain real value even facing take-it-or-leave-it offers.