Appendix to "Demand Inertia and the Hidden Impact of Pharmacy Benefit Managers"

Josh Feng, *University of Utah – David Eccles School of Business*
Luca Maini, *Harvard Medical School*

August 2023

A Additional Background on Pharmacy Benefit Managers

A.1 Evidence on concentration and prevalence of PBMs

The PBM industry is highly concentrated. As of 2019, the three largest PBMs (CVS-Aetna, Express Scripts, and Optum Rx) control approximately 80% of the market. During the time span covered by our analysis, the market was less concentrated, but not dramatically so. Exact data on market share is hard to come by, but Lipton et al. (1999) report that in 1997 the top four PBMs combined to reach approximately a 67% market share.[49] We also calculated a revenue-based measure of HHI using the PBM financial filings we describe in the main text and in the next section. The index fluctuates between 0.45 and 0.3 between 2000 and 2015 without a clear trend (Figure A1).

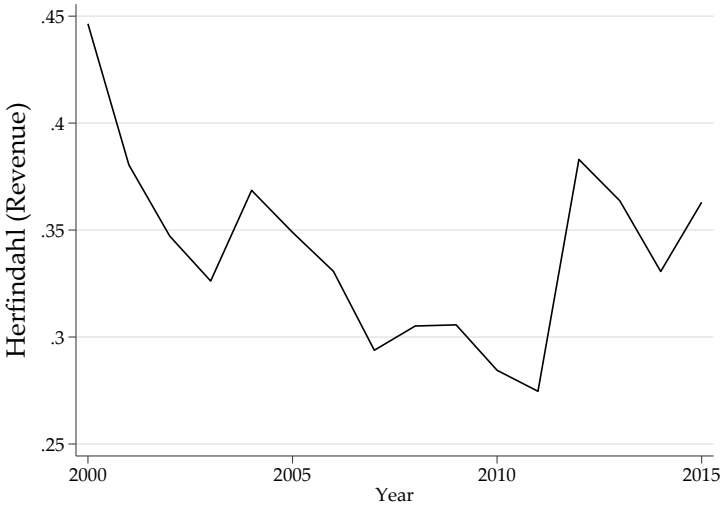

Figure A1: Revenue-based HHI in PBM market

[49] Authors' own calculations based on numbers in Table 1.

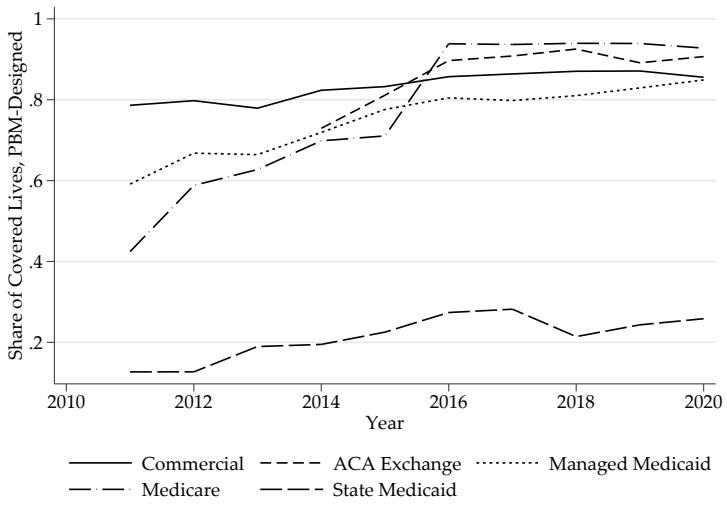

Figure A2: Fraction of lives with PBM-designed (custom or national) formularies

Notes: A plot of the share of lives in each insurance segment that is covered by a plan where the formulary is designed by a PBM, based on MMIT data.

The function of PBMs may differ slightly based on the market they serve. The majority of self-insured employers rely on PBMs to handle all aspects of prescription drug insurance, including negotiating prices, determining drug coverage through national formularies, and processing prescription drug claims. On the other hand, most health insurance companies only use PBMs for price negotiation, and design their own custom formularies. In the case of Medicare Part D, PBMs may also offer plans alongside traditional insurance providers. Lastly, some payers use PBMs solely for claim management, and make independent decisions regarding prices and coverage. Figure A2 shows the fraction of covered lives for whom PBMs negotiate drug prices across the commercial, Medicare, Health Exchanges, and Medicaid channels.

A.2 Description of the Bidding Structure

To guide our modeling decisions, we spoke with several former PBM industry executives and regulators and reviewed the recent Senate investigation report on insulin manufacturers.[50] In this section, we summarize what we learned.

The process of setting prices occurs separately for each drug class.[51] PBMs list several potential formulary arrangements to each drug manufacturer. Each formulary arrangement

[50] The report is available at https://www.finance.senate.gov/download/grassley-wyden-insulin-report, retrieved February 13, 2023.

[51] Cross-market agreements are possible but are not systematically used during the bidding process.

specifies the coverage tier for the drug owned by the manufacturer and for its therapeutic alternatives.[52] Manufacturers then send bids indicating the rebate they are willing to grant for each arrangement. In general, manufacturer surrender higher rebates for formularies that cover their drugs more generously and limit coverage for their competitors.

Based on these bids, PBMs create a set of "national" formularies, which they offer to payers. Payers can either choose a formulary offered by the PBM or create their own custom formulary by selecting arrangements from different drug classes. The price of the custom formulary is determined by the bids already submitted by each manufacturer.

PBMs face a trade-off in choosing national formularies, as they prefer lower prices but face costs when setting restrictive formularies. In particular, setting restrictions on popular drugs in chronic markets can result in payer dissatisfaction and may not lead to volume reduction due to consumer inertia and limited sensitivity to cost sharing.PBM Financial Filings

We use PBM financial filing information to estimate their gross profit margins.[53] PBM financial filings are complex. Some PBMs (e.g., Caremark) are divisions within larger companies. In addition, PBMs do not break down their revenues by drug market. However, they generally charge payers on a per member or per script basis. Finally, PBM revenues and margins are hard to measure, because the accounting measures incorporate payments for drugs that come from payers and are passed on to drug companies.[54]

We use the following process to build a list of PBMs and associated financial numbers:

1. We start off with a list of the largest PBMs in the 2011–2020 period based on MMIT data: Express Scripts, Medco, CVS/Caremark, Optum, Catamaran, Humana

2. For each year between 2000–2019 we look through the financial filings of each of these five PBMs, and check which competitors are mentioned. Using this approach, we add Catalyst (active 2000–2011) and NextRx (active 2000–2009). Others competitors mentioned include Prime Therapeutics and MedImpact, but neither is publicly traded, so we cannot collect information on them.

3. For each of the seven PBMs we collect revenue, gross profit, SG&A (which approximates overhead costs), net income, and operating income from financial filings

[52]The contracts we have seen do not mention specific competitors. However, they specify the number of other competing drugs in the same tier of coverage as the drug owned by the manufacturer in question. These contracts are found in the Appendices to the insulin report, which can be accessed through links available at the end of the report.

[53]The data used in our calculations is available online at https://sites.google.com/view/jfeng/data.

[54]A 2017 article from Bloomberg provides a more detailed discussion of this issue: bloomberg.com/news/articles/2017-03-31/drug-middlemen-have-slim-profit-margins-just-ask-them, retrieved February 17, 2023.

4. We then compute variable cost as the difference between revenue and gross profit

The revenues number include both rebates from drug companies and fees payed by clients, whereas profit is what the PBM retains. Hence, variable cost should primarily reflects payments to drug companies, plus some small processing costs for prescription management. If we assume that other costs are small enough, we can approximate the PBMs gross profit margin (i.e., not including fixed costs) as

$$\text{PBM Margin} = \frac{\text{PBM Gross Profit}}{\text{PBM Revenues} - \text{PBM Gross Profit}} = \frac{\text{PBM Revenues} - \text{PBM Variable Cost}}{\text{PBM Variable Cost}}$$

Our estimated margin is approximately 6.5 percent, and is fairly consistent over time for the largest PBMs. This margin comes in part from retaining a fraction of the rebates, and in part from fees charged to payers. Express Scripts claims that it keeps about 10 percent of all rebates and passes the rest onto payers.[55] If we normalize list prices to \$1 and assume an average discount of 30 percent, Express Scripts would keep \$0.03 of the \$0.30 discount, which would equal 4.3 percent of variable cost (\$0.70).

B Additional Reduced-Form Evidence

B.1 Alternative Market Definitions

We present evidence that the empirical facts documented in Section 2 are robust to definition of drug class. The results in the main text use ATC-4. Drugs with the same ATC-4 generally have very similar chemical properties and treat the same disease. ATC-4 is commonly used in the pharmaceutical literature to define classes (see, e.g., Dubois et al., 2019).

Here, we show results using three alternative class measures. The first is the ATC-3, which is the next higher tier in the ATC classification. The next two are proprietary classifications that are included in our MMIT and SSR Health data.

Correlation Between Net Price and Formulary Coverage We run the cross-sectional regressions using the three classifications not reported in the main text. Table A1 reports the results. We consistently find that net price has an impact on formulary coverage generosity, particularly the rate of preferred coverage. List price, on the other hand, has a noisily-

[55] Wall Street Journal article, "Drugmakers Point Finger at Middlemen for Rising Drug Prices," available at https://www.wsj.com/articles/drugmakers-point-finger-at-middlemen-for-rising-drug-prices-1475443336, retrieved February 17, 2023.

measured correlation with formulary generosity in all specifications without any consistency in the magnitude or sign of the coefficient.

Table A1: Prices vs. Formulary Correlations

(a) ATC-3

	Frac. Preferred		Frac. Covered	
Log(net price)	-0.0323**	-0.0344*	-0.00888**	-0.00868
	(0.00616)	(0.0134)	(0.00246)	(0.00625)
Log(WAC)		0.00303		-0.000293
		(0.0150)		(0.00742)
Fixed Effects	Class-by-Year			
Observations	2520	2520	2520	2520

(b) MMIT Class

	Frac. Preferred		Frac. Covered	
Log(net price)	-0.0329*	-0.0582*	-0.0126*	-0.0233*
	(0.0140)	(0.0245)	(0.00631)	(0.0107)
Log(WAC)		0.0397		0.0168
		(0.0348)		(0.0119)
Fixed Effects	Class-by-Year			
Observations	2740	2740	2740	2740

(c) SSR Health Class

	Frac. Preferred		Frac. Covered	
Log(net price)	-0.0226**	-0.0286**	-0.00837**	-0.00861
	(0.00514)	(0.00857)	(0.00296)	(0.00607)
Log(WAC)		0.00889		0.000362
		(0.0119)		(0.00665)
Fixed Effects	Class-by-Year			
Observations	2425	2425	2425	2425

Notes: Estimates of Equation 1, comparing within class-by-year cells. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Response to Branded and Generic Competitor Entry Next, we present results on response to entry in Table A2. We report results from a simplified version of Equation 2

$$y_{it} = \alpha_i + \gamma_t + \beta_1 \times BrandedEntryCounter_{it} + \beta_2 \times GenericEntryCounter_{it} + \varepsilon_{it} \quad \text{(A1)}$$

We consistently find across all the class definitions that entry events have a minimal effect on net prices. There is weak evidence in some cases of a small drop in net price in response to competitor loss-of-exclusivity (generic versions of competitors enter), but the estimates rule out any large effects. The only consistent and significant result we find is a units response to competitor loss-of-exclusivity. The units sold drops somewhere between 10–20 percent for the incumbent branded drug.

Table A2: Response to Competitor Entry

(a) ATC-4

	log(WAC)	log(net)	log(units)	Preferred	Covered
Competitor	-0.047**	0.007	0.097**	0.002	0.013**
Launch Counter	(0.015)	(0.023)	(0.035)	(0.004)	(0.003)
Competitor	0.0079	-0.014	-0.118+	-0.007	0.003
LOE Counter	(0.023)	(0.040)	(0.069)	(0.010)	(0.005)
Drug FE	x	x	x	x	x
Year FE	x	x		x	x
Age FE			x		
Observations	6706	5392	6701	2941	2941

(b) ATC-3

	log(WAC)	log(net)	log(units)	Preferred	Covered
Competitor	-0.044**	0.016	0.094**	0.003	0.009**
Launch Counter	(0.011)	(0.018)	(0.030)	(0.004)	(0.002)
Competitor	0.0082	-0.041	-0.110**	-0.001	-0.001
LOE Counter	(0.023)	(0.025)	(0.035)	(0.006)	(0.003)
Drug FE	x	x	x	x	x
Year FE	x	x		x	x
Age FE			x		
Observations	6706	5392	6701	2941	2941

(c) MMIT Class

	log(WAC)	log(net)	log(units)	Preferred	Covered
Competitor	0.005	0.001	0.093	-0.003	-0.003
Launch Counter	(0.028)	(0.041)	(0.080)	(0.012)	(0.007)
Competitor	0.030	0.019	-0.448**	-0.015	-0.012*
LOE Counter	(0.031)	(0.039)	(0.088)	(0.012)	(0.005)
Drug FE	x	x	x	x	x
Year FE	x	x		x	x
Age FE			x		
Observations	7061	5796	7055	3372	3372

(d) SSR Health Class

	log(WAC)	log(net)	log(units)	Preferred	Covered
Competitor	-0.027*	0.008	-0.006	-0.0002	0.008**
Launch Counter	(0.012)	(0.019)	(0.033)	(0.004)	(0.003)
Competitor	0.053**	0.035	-0.149**	0.006	-0.002
LOE Counter	(0.017)	(0.034)	(0.052)	(0.007)	(0.004)
Drug FE	x	x	x	x	x
Year FE	x	x		x	x
Age FE			x		
Observations	5883	5211	5881	2609	2609

Notes: Estimates of Equation A1, using different class definitions to identify competitors. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table A3: Prices vs. Formulary Correlations with Manufacturer Fixed Effects

	Frac. Preferred		Frac. Covered	
Log(net price)	-0.0266**	-0.0512**	-0.0127	-0.0230
	(0.00666)	(0.0149)	(0.00806)	(0.0160)
Log(WAC)		0.0322+		0.0136
		(0.0169)		(0.0184)
Fixed Effects	Class-by-Year, Manufacturer			
Observations	2520	2520	2520	2520

Notes: Estimates of Equation 1, comparing within class-by-year cells. The results here include a manufacturer fixed effect. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

B.2 Accounting for Manufacturer Fixed Effects

We provide additional specifications related to the cross-sectional tradeoff within class between net prices and formulary coverage. As noted in the main text, our evidence should not be interpreted causally. Here, we included manufacturer fixed effects to try to account for systematic differences across companies in drug quality and negotiation experience, which they could apply across all of their drugs in different disease areas. Table A3 presents the estimates.

B.3 Suggestive Evidence of the Strategic Impact of Inertia

To provide suggestive evidence on how inertia affects incumbents' response to entry, we run a regression where we allow the price response to vary linearly by incumbent age at time of competitor entry. The idea is that incumbents that have been on the market for longer will have a larger patient base, and therefore are less likely respond to competitor entry by lowering prices and trying to obtain better formulary coverage. Formally, we run the following specification

$$y_{it} = \alpha_i + \gamma_t + \beta_1 \times BrandedEntryCounter_{it} + \beta_2 \times BrandedEntryCounter_{it} \times Age_i + \varepsilon_{it} \quad \text{(A2)}$$

The first term allows for a response size of $\beta_1$at age zero. The second term, computed by incrementing the counter by the number of years an incumbent has been on the market ("age") at the time of competitor entry, allows for the response size to vary linearly with the incumbent's age at the time of entry. Table A4 reports the results. Confirming our intuition,

we find a negative, albeit statistically insignificant, net price response at age zero of about six percent. This negative effect is then offset by a positive interaction term. The coefficients combine to suggest that the negative price response goes away by age seven. We find the opposite result for preferred coverage. Younger incumbents respond to entry by obtaining more favorable formulary position. This is then offset by a negative interaction term, so older incumbents appear to have weaker incentives to obtain better formulary position in response to competitor entry.

Table A4: Incumbent Response Variation by Age at Entry

	log(net)		Preferred	
Competitor Launch Counter	0.00614	-0.0599	0.00143	0.0319**
	(0.0242)	(0.0431)	(0.00339)	(0.00664)
Age Scaled Competitor Launch Counter		0.00833+		-0.00363**
		(0.00473)		(0.000721)
Drug FE	x	x	x	x
Year FE	x	x	x	x
Observations	5392	5392	2941	2941

Notes: Estimates of Equation A2. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

B.4 Hepatitis C Case Study

Finally, we present evidence from SSR data showcasing price trends in the Hepatitis C market. In this market, demand inertia is not relevant thanks to the recent introduction of novel therapies that can cure the disease in a few weeks. The first cure for Hepatitis C, Sovaldi, entered the market in 2013, and was quickly followed by a series of competitors. Figure A3 shows that net prices have declined each year since 2013.[56]

C Data Construction and Structural Estimation Details

C.1 Construction of pre-2007 net price series

Our structural analysis covers 1996–2013. Since SSR Health only tracks net prices starting in 2007, we collect additional data on US net revenues and volume sold by drug and year

[56]The declining prices are exacerbated by the presence of a different kind of dynamic effect: market shrinking due to the exhaustion of the stock of patients. Prior to 2013, Hepatitis C treatments were much less effective and had significant side effects. As a result, many patients simply chose to remain untreated. When the new cures arrived on the market, the stock of potential patients was very large, but it quickly shrank because the number of treated patients was much greater than the flow of new patients. This environment likely exacerbated price competition as firms raced to treat as many patients as possible.

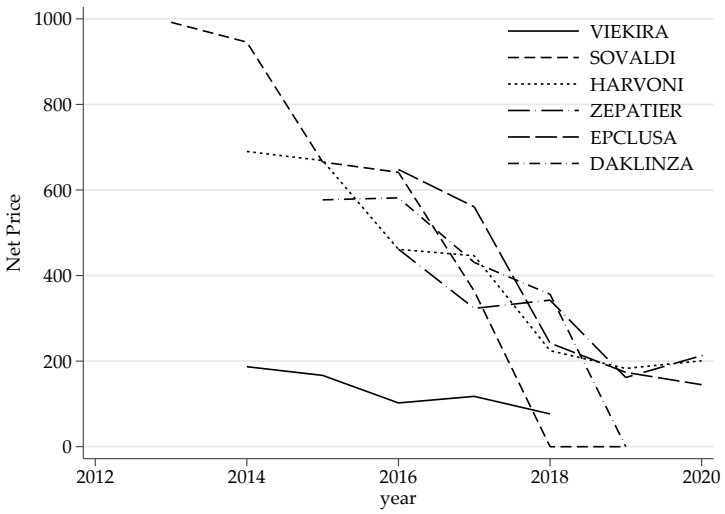

Figure A3: Hepatitis C Drug Prices

going back to 1996 to arrive at average negotiated prices. This requires some additional assumptions, which we detail in this section.

Calculation of US volume sold We start by calculating volume sold. We use MEPS, which is available throughout the entire 1996–2013 span. We begin by collecting scripts data at the drug-dosage level in each year:

1. We extract all scripts in MEPS related to Zocor, Crestor, Lipitor, Mevacor, and Vytorin.[57]

2. We take the RXQUANTY variable (days supplied) and multiply it by the weight given to each respondent

3. We aggregate total quantity by product, strength, and year

We break out quantity by dosage, because we want to focus on the medium intensity market. Table A5 summarizes the dosages and treatment intensities associated with different medications, along with total quantity of brand sales (number of pills) from 1996–2013 calculated based on MEPS data. The medium intensity market is by far the largest, and the majority of sales of Zocor, Lipitor, and Crestor are in the medium intensity segment. This supports our modeling assumption that they are mainly competing against each other rather than against Vytorin and Zetia.

[57] We only include Mevacor because Merck reports aggregate sales for Mevacor and Zocor in the early years of our sample.

Table A5: Drug and Dosage vs. Treatment Intensity

Drug	Dosage	Treatment Intensity	Number of Units (billion)	Share Within Molecule	Share within the moderate intensity market
Lipitor	10 MG	Moderate	15.6	0.47	39.0%
Lipitor	20 MG	Moderate	10.8	0.32	27.0%
Lipitor	40 MG	High	5.2	0.16	
Lipitor	80 MG	High	1.7	0.05	
Vytorin	10 MG-10 MG	Moderate	0.3	0.07	0.8%
Vytorin	10 MG-20 MG	High	1.9	0.42	
Vytorin	10 MG-40 MG	High	1.7	0.38	
Vytorin	10 MG-80 MG	High	0.6	0.13	
Crestor	5 MG	Moderate	1.2	0.14	3.0%
Crestor	10 MG	Moderate	4.2	0.49	10.5%
Crestor	20 MG	High	2.4	0.28	
Crestor	40 MG	High	0.8	0.09	
Zocor	5 MG	Low	0.5	0.04	
Zocor	10 MG	Low	2.1	0.18	
Zocor	20 MG	Moderate	4.9	0.42	12.3%
Zocor	40 MG	Moderate	3.0	0.25	7.5%
Zocor	80 MG	High	1.3	0.11	

Notes: Categorization of different strengths of each drug into treatment classes and the associated units sold for the branded version of each dosage from 1996–2013. The focus of the analysis is on Zocor 20mg, Lipitor 10mg, and Crestor 10mg, by far the most popular dosages for starting patients. Vytorin strengths report Zetia first, Zocor second. 10 MG-10MG Vytorin has negligible market share. Sources: Prescriber's Letters (2011), available at: https://www.epocrates.com/sites/default/files/res/dacc/2017/StatinDoseTablePL1703.pdf, retrieved July 15th, 2021.

Calculation of net US revenue Drug manufacturers report data as follows over the 1996–2007 period (all data comes from 10-K filings unless otherwise noted):

1. 1996–1999: Merck reports global net sales for all its cholesterol drugs (Mevacor and Zocor); Pfizer reports Lipitor global net sales

2. 2000–2003: Merck reports Zocor global net sales separately; Pfizer reports Lipitor global net sales

3. 2004–2006: Merck reports US net sales for Zocor in its conference calls

4. 2003–2013: all companies report net US sales in its 10-K filings.

To supplement the data above, we use IMS Top-Line Market Data, which tracks US and global gross revenues by drug, and is available starting in 2000.[58]

Using these data sources, we use the following procedure to calculate net US revenues:

1. In years with reported net U.S. revenue, we just use the number available.

2. In years where only net global revenue are reported, we arrive at a net U.S. revenue number by making the assumption that invoice sales are equal to net sales in non-US markets. Our exact formula is

$$\begin{aligned}\text{Net US Revenue} = &\text{Net Global Revenue} \\ &- (\text{Gross Global Revenue, IMS} - \text{Gross US Revenue, IMS})\end{aligned}$$

3. We adjust Zocor sales in 2006 to account for the fact that generics entered in the middle of the year. We use quarterly revenue from the first two quarters and project out to the whole year.

4. In the 1996–1999 period, we multiply world net sales by the U.S. to world net sales ratio from 2000 to arrive at an estimate for U.S. net sales. We further multiply U.S. net sales reported by Merck for all anti-cholesterol drugs by the ratio of Zocor scripts to total scripts (Zocor plus Mevacor) to arrive at a net sales number for Zocor.

To validate the assumption in (2) we plot the difference in US gross-to-net sales and Global gross-to-net sales in Figure A4. The figure shows that the two quantities are roughly equivalent.

Finally, we collect data on gross prices. To do so, we collect average wholesale price (AWP) and median price paid (labeled as "pay" in the MarketScan data) for each drug by dosage by seller from the MarketScan data. Each entry in the data contains these two quantities. The AWP is constant at the national level at a given point in time, while the "pay" variable partially reflects pharmacy markups.

Construction of Net Price Series Under Alternative Assumptions A major challenge in terms of calculating net prices is that net sales estimates are at the drug level, but different dosages of the same drug are sometimes priced differently. Otherwise, we could just divide net sales by quantity. For example, a 40MG Lipitor tablet has a gross price that is

[58] Based on business news articles and Top-Line Market Data: http://www.imshealth.com/en/about-us/news/top-line-market-data, retrieved January 14, 2017. The Top-Line data is available on Wayback Machine, and we have also included all IMS revenue data as an additional tab in the PBM financials excel file.

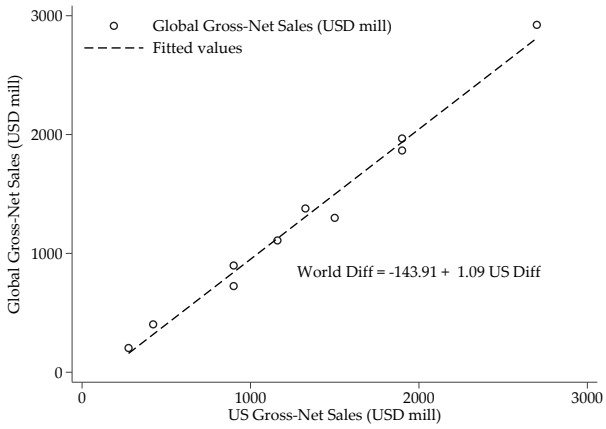

Figure A4: Global vs. US Rebates ($bill)

Notes: A binscatter showing a comparison of discounts and rebates between the US market and global numbers, covering the three main statins in the years 2007–2011. The calculations are based on US and global gross revenue numbers from IMS Topline and company-reported US and global net revenue.

$1 more expensive than the more commonly used 10MG tablet. Crestor, on the other hand, sets the same price for all dosage levels. Without additional knowledge, we cannot pin down the negotiated price for medium intensity dosages of each drug, which is the focus of our analysis.

To solve this problem, we assume that percentage discounts per pill are equal across dosages for the same drug. To do so, we compare the difference in overall gross and net revenue for each drug. We then take the percentage discount off the list price of the medium intensity dosage. In cases where there are multiple dosages and different prices associated with each, we take a weighted average by volume. In the end, we arrive at average negotiated price in the U.S. market for each of the three drugs in each year.

We also compute two additional prices series. Figure A5 presents a comparison of the results from the three approaches.

As a first alternative, we derive net prices assuming constant dollar discounts across dosages. In the core derivation, the percentage discount per pill is assumed to be the same across dosages. An alternative assumption is that dollar discounts are the same across dosages. There are minimal differences between the two approaches. The biggest deviations are about 5 cents, and Crestor is barely affected, as its dosages have similar prices. The differences that are present depend on how the prices of low-intensity and high-intensity dosages relate to those of the medium-intensity dosages. When high-intensity dosages are priced higher, they absorb more of the discount under the percentage model, which results in higher prices.

Figure A5: Price Comparison – Three Approaches

(a) Zocor

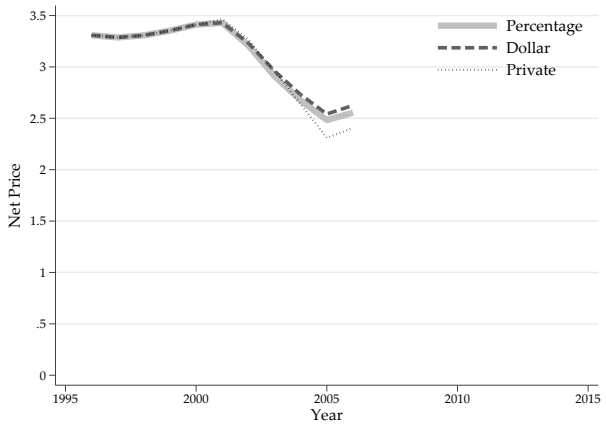

(b) Lipitor

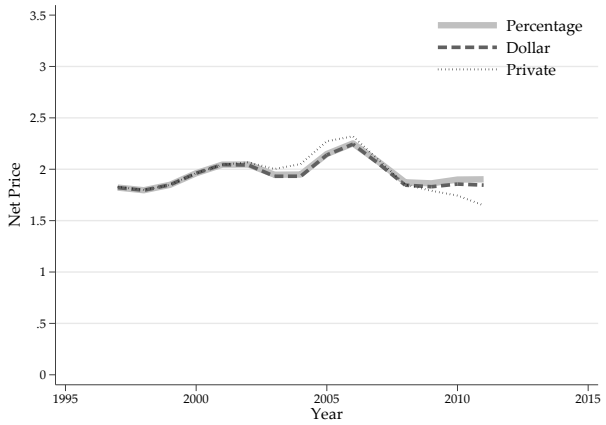

(c) Crestor

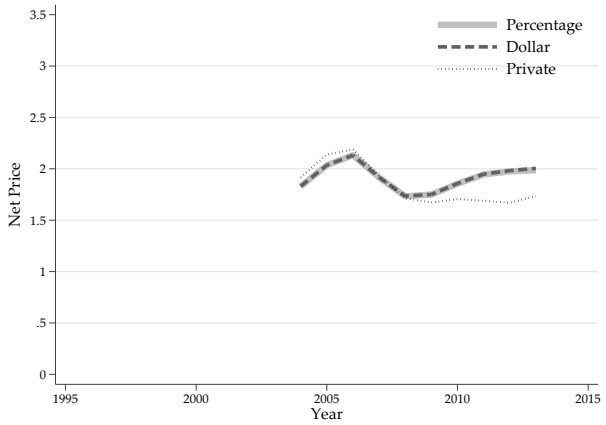

Notes: Net prices under the three different methodological approaches.

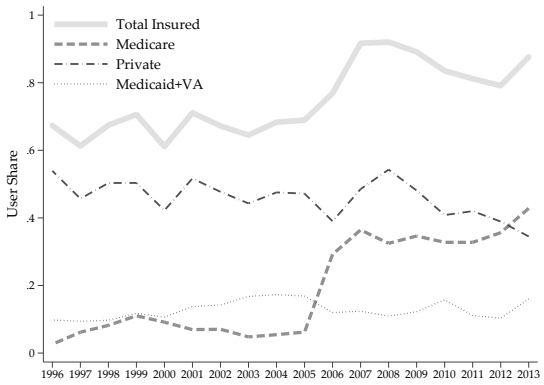

Figure A6: Insurance Coverage Type for Cholesterol Patients

Notes: A plot of the percentage of patients taking anti-cholesterol medication who are covered by each type of insurance (a small percentage are covered by multiple). Calculations based on MEPS data.

A more substantive adjustment is to isolate the negotiated prices in markets covered by PBMs, which generally results in lower derived prices. PBMs serve Part D plans and private insurance, which does not include Medicaid, Veterans Administration (VA), and uninsured patients. Uninsured patients pay an amount close to list price, and make up about 15–30 percent of the individuals taking anti-cholesterol medication. Medicaid and VA patients pay less than or equal to the private market, and account for 10–15 percent of people. Figure A6 documents these numbers. Our approach is to first purge the net revenue data of uninsured patients paying list price, which pushes down the effective net price in private markets. Then, to create as high a potential negotiated price as possible, we assume that the VA pays \$0 and that Medicaid receives a 30 percent discount off the net price in the private market. We then obtain the net price by solving out the following accounting identity for p_{PBM}:

$$p_{PBM}(s_{Medicare} + s_{Private} + 0.7s_{Medicaid}) + s_{uninsured}p_{list} = p_{net}$$

where s represents the share of individuals taking medication covered under each type of insurance, p_{net} is the overall net price derived earlier, and p_{list} is the list price. Even this adjustment does not lead to substantial changes in the price series, with the exception of the price of Crestor in 2009—2013. Column 3 of Table 4 reports estimates that use this price series. The main difference is that the penalty of excluding a drug altogether is lower in this specification, but the overall implications of the model are quite similar. In our counterfactuals we stick with the constant percentage discount price series because it requires fewer assumptions.

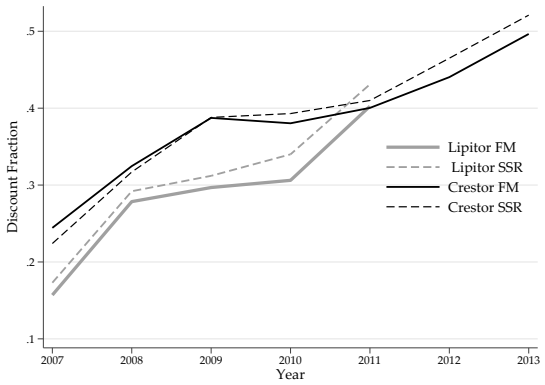

Figure A7: Comparison of Percentage Discounts

Notes: FM refers to our discount series calculated using the methodology described in this Appendix. SSR refers to percentage discounts calculated by SSR Health using proprietary data.

Comparisons to SSR Health Estimates Figure A7 provides a comparison of the percentage discounts calculated using our approach versus the SSR Health data. Our methodology approximates the SSR Health data very accurately for both Lipitor and Crestor.

C.2 Computational Step for the Estimation of the Bidding Model

Matching estimated net prices series to the model's prediction is the core exercise of our paper. In this section we describe how we perform the estimation algorithm:

1. Start from a guess for the parameters in the PBM's objective function: $\theta' = \left(\theta'_s, \theta'_e, \theta'_j, \theta'_\omega\right)$
2. Compute the value function for each firm, starting from the last period—when only Crestor is active—and solving the model using backwards induction:
   (a) In period t, fix a value x_t for the state space, and compute optimal prices and payoffs for each firm as follows:
      i. Guess an initial vector of prices $(p^0_{ZOCOR}, p^0_{LIPITOR}, p^0_{CRESTOR})$
      ii. For each drug j, starting with Zocor and ending with Crestor, find $p^\star_j\left(p^0_{-j}\right) = \arg V_{jt}\left(x_t\right)$, where $V_{jt}\left(x_t\right)$ is the value function of drug j defined as the maximization problem in equation 10. Intuitively, $p^\star_j\left(p'_{-j}\right)$ is the best response of manufacturer j holding fixed the prices of the other two drugs
      iii. Replace p^0_j with $p^1_j = p^\star_j\left(p^0_{-j}\right)$ in the initial vector of guesses. Repeat step ii. for all drugs that have not lost exclusivity yet (prices for generics are fixed).
      iv. Repeat steps ii. and iii. until $\sum_j \left(p^n_j - p^{n-1}_{-j}\right)^2 < \varepsilon$ for some pre-established threshold ε (we use $\varepsilon = 10^{-8}$). This is the equilibrium.

v. Calculate the value function for all drugs j at the estimated values $(p^n_{ZOCOR}, p^n_{LIPITOR}, p^n_{CRESTOR})$

(b) Repeat step (a) for all values $x_t \in [0, 0.05, 0.12, 0.2, 0.3, 0.5, 0.7]^3$ (excluding combinations where the sum of the market shares sums to more than one).

(c) To find the value function for values of x_t that are not in $[0, 0.05, 0.12, 0.2, 0.3, 0.5, 0.7]^3$ we use the scattered interpolant function in MATLAB, with the natural neighbor option, which yields a smoother function than a simple linear interpolation.[59]

(d) Move to period $t-1$, and repeat steps (a)–(c)

3. Once we have a value function for each period and state, we play out the game, giving Zocor a 0.4 market share coming into 1996. This yields a set of predicted equilibrium prices associated with the parameter vector θ', which we can then compare to the observed net prices in the data.

Given the routine for computing equilibrium prices given a parameter vector, we then run a search for the parameter vector that minimizes the distance between predicted and observed net prices. The routine is computationally intensive. Each evaluation takes 5–10 minutes using parallelized code on a cluster with 15 workers. The main bottleneck is the set of periods during which there are three active players. In these periods, the state space is very large.

C.3 Evidence of uniqueness of equilibrium

The algorithm described above is based on a minimum distance estimator, which is only valid if the game has a unique equilibrium. While we do not have a formal proof of uniqueness, we offer some heuristic analysis that suggest multiplicity is not an issue in our model.

Because our game has a finite horizon, a sufficient condition for uniqueness is that each stage game has a unique equilibrium. Results in the literature have provided conditions for existence uniqueness in price setting games under certain classes of demand systems, focusing on price and market share bounds (Thomadsen, 2005; Aksoy-Pierson et al., 2013). While our demand system fits under the conditions, our model also contains an intermediary, which will alter the effective demand curve facing firms. This makes it difficult to directly apply the conditions derived in the existing literature.

To show that our stage games are likely to have a unique equilibrium, we calculate the best response functions in the years during which only two manufacturers compete (1997–2003 for Zocor and Lipitor, and 2007–2011 for Lipitor and Crestor). This allows for a simple

[59] We use a sparse vector to reduce computational time. The vector is more densely populated at lower market shares because those states are more likely to be observed in the data.

view of best response functions in a two-player game. While the best response functions are non-monotonic, we do find that they only intersect once for the values of the parameter θ that we estimate and for the values of the state space that we observe. Figure A8 plots the best response functions for the five years in which Lipitor and Crestor competed. One thing to notice is that prices are not particularly responsive to one another. This is largely due to the fact that generic Zocor has a large market share, so undercutting the branded competitor has a relatively small impact on market share. We conducted additional simulations (not shown) confirming that if generic Zocor had lower market shares, the best response functions would be more responsive. In these simulations too, we find that the functions satisfy the single-crossing property.

D Additional Structural Model Discussion and Results

D.1 Copay elasticities and first-stage estimates

One potential issue in the demand estimation is that we do not have a suitable instrument for copays. While instrumenting for copays is desirable, it is not clear what bias the absence of an instrument introduces. Patients could select into plans that have a low copay associated with the drugs they like, leading to an upward bias in the estimated copay elasticity. On the other hand, copays might be higher for higher quality drugs, or drugs with more loyal patients because manufacturers might be more willing to accept worse coverage in exchange for charging higher net prices. Hence, our estimates might be biased upward or downward, and, in practice, these two effects might cancel each other out.

To validate our estimates, we compare them to the estimates from Einav et al. (2016), which exploits the Medicare Part D donut hole as an exogenous source of variation in copays. Einav et al. (2016) use data from 2007–2011, and estimate copay elasticities for the most popular branded drugs, including Lipitor and Crestor. They estimate that Lipitor and Crestor have a copay elasticity of -0.33 and -0.31 respectively.

When we compute the copay elasticities of Lipitor and Crestor using our demand system in the same 2007–2011 period, we find that our demand system generates similar implied elasticities. Figure A9 plots the elasticities for Lipitor and Crestor by year. Our estimates for Lipitor range from -0.26 to -0.42 and our estimates for Crestor range from -0.43 to -0.51. These are quantitatively similar to the findings in Einav et al. (2016). The differences likely arise from our imposition of a functional form and differences in sample. Implicitly, Einav et al. (2016) are measuring the response of all patients who reach the donut hole, potentially a slightly selected sample of older individuals who take more drugs. These patients are likely

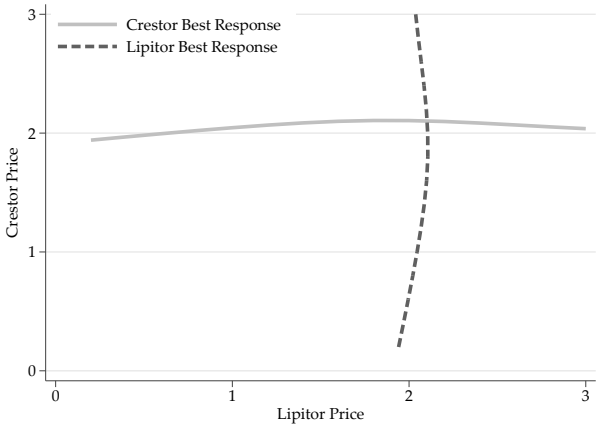

(a) 2007

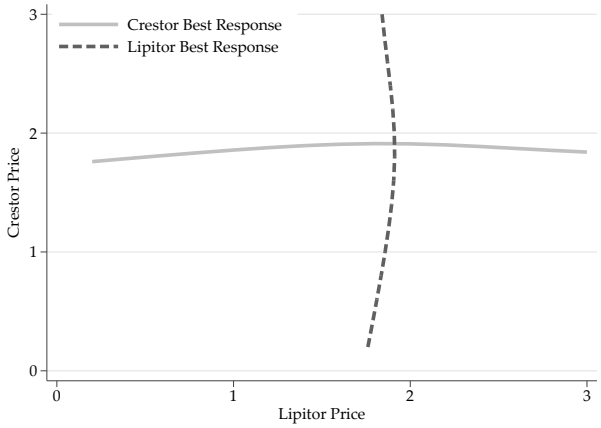

(b) 2008

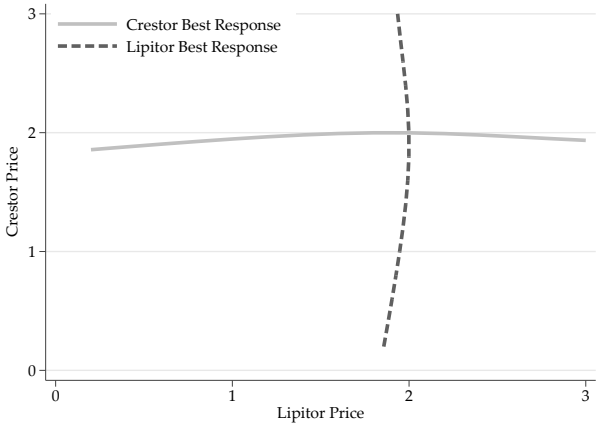

(c) 2009

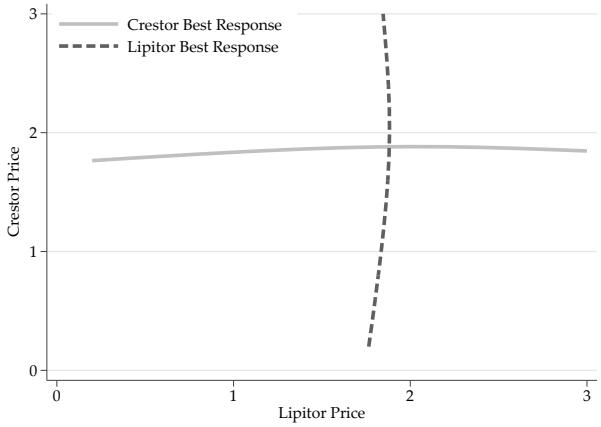

(d) 2010

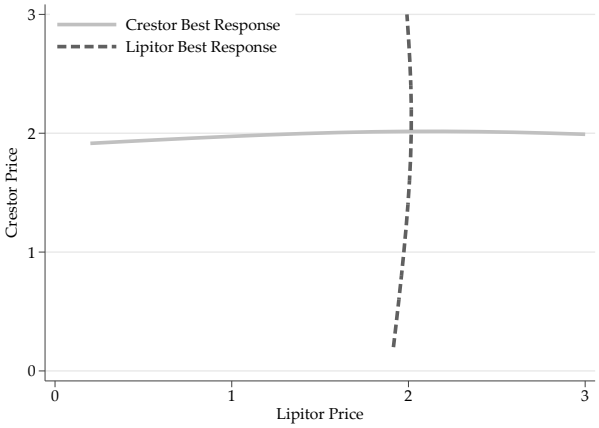

(e) 2011

Figure A8: Best Response Function Plots

Notes: A plot of best response functions for Lipitor and Crestor after the entry of generic Zocor and before the entry of generic Lipitor (2007–2011). In each year, we show the state that is closest to the observed state in equilibrium play: 2007, [0.5, 0.3, 0.05]; 2008, [0.5, 0.2, 0.05]; 2009, [0.5; 0.2; 0.05]; 2010, [0.7; 0.12; 0.05]; 2011, [0.7;0.12;0.05].

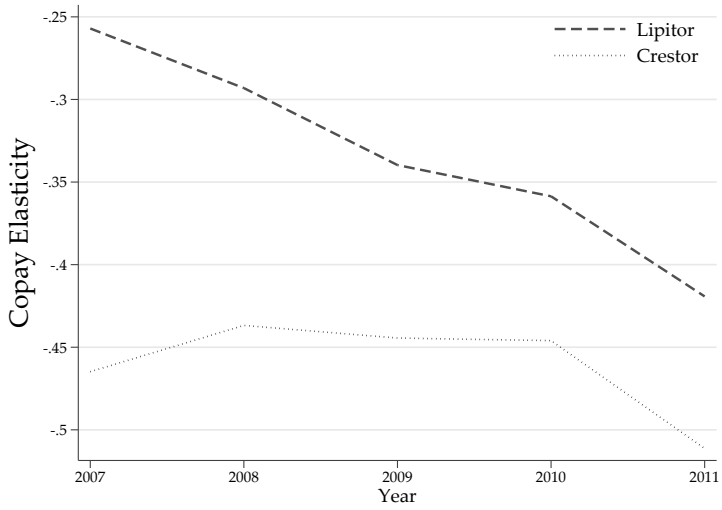

Figure A9: Predicted Copay Elasticities 2007–2011

Table A6: Distribution of Copay Changes

	5th	10th	25th	Median	75th	90th	95th
Change in Monthly Copay	-$10	-$2	$0	$0	$0	$5	$11

Notes: Distribution of year-to-year copay changes at the person-drug level, using the MarketScan sample used in demand estimation.

to be more inelastic than the overall population.

In addition, we provide statistics in Table A6 on the distribution of year-to-year copay changes at the person-drug level in our MarketScan sample. Copay changes only occur in 44 percent of person-years. For reference, the median copay for branded statins is $15 per month and the average is $21.70.

Finally, we report the first stage regressions associated with the instrumental variables approach, which is used in both the 2SPS and 2SRI estimation approaches. Table A7 reports the results. As shown in the table, the leave-one-out cohort mean is strongly predictive of a given indivdual's previous choice indicators. The coefficients are close to one, which makes sense given that we are predicting shares across all individual-drug pairs, and each cohort has a large number of individuals.

Table A7: First Stage Instrumental Variables Estimates

	2SPS and 2SRI	
	Incumbent	Non-Outside
Leave-one-out measure (Incumbent)	0.99988	
	(0.00024)	
Leave-one-out measure (Non-Outside)		0.99981
		(0.00009)
R^2	0.2564	0.7311
F	>99,999	>99,999
N	49,696,288	49,696,292

Notes: Incumbent refers to the term $I_{s_{it-1}=j,j\neq 0}$ and non-outside refers to the term $I_{s_{it-1}\neq 0,j\neq 0}$. We construct the cohort means ("leave-one-out measure") based on Equation 11, and run first-stage regressions based on Equation 12.

Table A8: Parameter estimates using different "offline" parameters

	Estimates			
Parameter	Baseline	2SRI	$\beta = 0.91$	$\beta = 0.85$
θ_s	1.70	1.74	1.66	1.65
θ_e	1.50	1.39	1.44	1.48
θ_j	1.11	0.86	0.94	0.94
θ_ω	0.063	0.050	0.055	0.064
Total error	0.46	0.43	0.48	0.47

Notes: All parameter estimates are based on the "Pct Discount" price series. Total error refers to the sum of squared errors across all branded drug-years between 1996 to 2013 (excluding periods after generic entry for each drug).

D.2 Estimation results under different demand parameters and discount rates

For our core specification, we make several choices related to the offline parameters that are then inputted into the estimation of the dynamic game. This includes which demand parameters to use (2SPS in the core results) and picking a discount factor for manufacturers (0.88 in the core results). Table A8 provides estimates using 2SRI demand parameters and using different discount rates. The model fit is similar across the specifications, and the parameter estimates are also similar.

D.3 Formulary Predictions

The dynamic estimation of our model targets the net prices observed in the data. However, our model can also generate predictions for formulary coverage, given those prices. Figure A10 presents the predictions of the estimated model for formulary coverage of each drug. The model creates a probabilistic distribution over all possible formulary arrangements. For each drug, we plot the average expected copay facing patients (conditional on formulary inclusion) and the probability the drug is excluded altogether.

We find that our model's predictions match qualitative and quantitative evidence on formulary structure in the statin market. The model predicts that Zocor is more likely to be placed in the non-preferred tier in earlier years, in line with both anecdotal evidence and the mean daily copay series presented in Figure 3 of the main text.[60] We note that copays are increasing for all drugs over time. This does not come from the model itself—which predicts formulary tier—but from the MarketScan data, which shows increasing copays at all tier levels over time.

The model also predicts low exclusion rates in earlier years, with some exclusion for Crestor.[61] Figure A11 plots the formulary exclusion rates for Lipitor and Crestor in Medicare Part D formularies starting in 2008, when data becomes available. The two drugs also face softer restrictions such as step therapy and prior authorization, outcomes not included in the model, which make a direct quantitative comparison infeasible.

Finally, we also plot the probability that multiple drugs are excluded in equilibrium. As discussed in the counterfactual analysis in Section 5.1, if PBMs cannot exclude more than one drug, the equilibrium prices and spending are significantly higher. However, Figure A12 shows that this is a rare outcome in equilibrium—below 2 percent in all years except the years before generic entry.

[60] For anecdotal evidence, see Table 9 of a 2001 report by the California HealthCare Foundation. The table shows that all the major PBMs put Lipitor in the preferred tier, whereas Zocor was only in the preferred tier of the formularies of Medco and AdvancePCS: "Prescription Drug Coverage and Formulary Use in California: Difference Approaches and Emerging Trends" available at: http://www.chcf.org/publications/2001/06/prescription-drug-coverage-and-formulary-use-in-california-different-approaches-and-emerging-trends. Note that the MEPS series is plotting copays conditional on choosing a drug, whereas our model is predicting the unconditional average copay facing patients.

[61] Crestor initially was excluded from some formularies. See The Guardian article "US blow to AstraZeneca": https://www.theguardian.com/business/2003/oct/04/4.

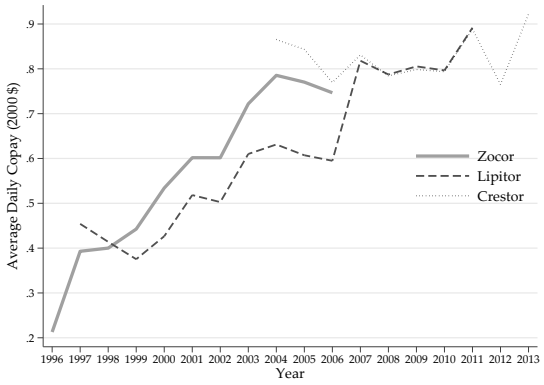

(a) Average Copay (Conditional on Inclusion)

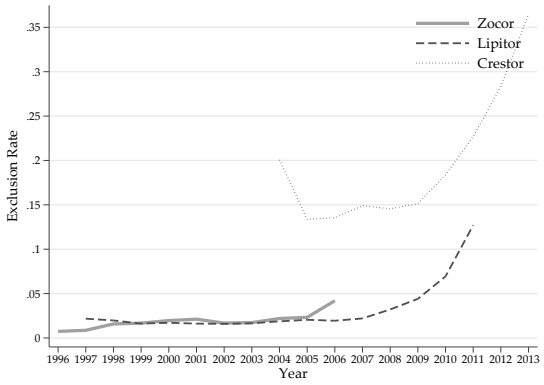

(b) Exclusion Rate

Figure A10: Additional Model Predictions – Formulary Outcomes

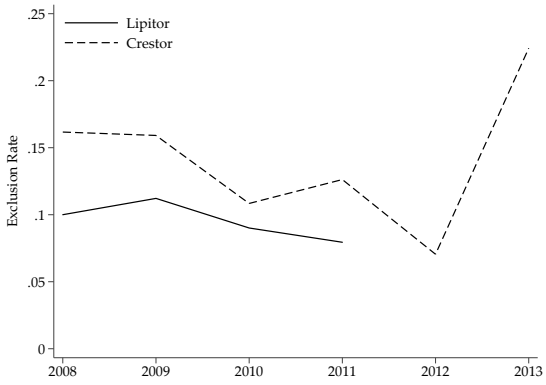

Figure A11: Formulary Exclusions for Anti-Cholesterol Drugs (Part D)

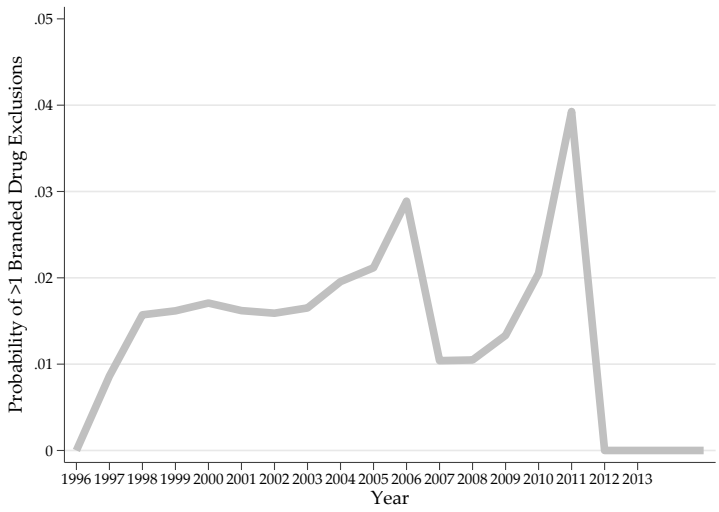

Figure A12: Double Exclusion Rates in Equilibrium

Notes: A plot of the rate at which the representative PBM excludes multiple brand drugs from the formulary, based on predicted play under the core model estimates.

D.4 Estimation results under alternative specifications of the PBM's objective function

We use a flexible control function approach to model the PBM's objective function. To test the robustness of this approach, we re-estimate the model using three alternative specifications. The results are qualitatively identical. We describe our additional analysis here.

Allowing for dynamic considerations Just like firms are forward-looking, PBMs could also incorporate dynamic considerations in their decision making. For example, PBMs might prefer starting new patients onto drugs who are about to experience loss-of-exclusivity, in order to reduce their future costs. At the same time, since patients frequently switch insurance, PBMs may only realize a small fraction of these potential future savings. We test for PBM dynamic incentives by introducing an additional term:

$$\beta^{PBM} V_{t+1}^{PBM}\left(x_{t+1}\right)$$

The term $V_{t+1}^{PBM}\left(x_{t+1}\right)$ represents the PBM's value function, and it's calculated dynamically through the same procedure used to compute the manufacturers' value functions we describe in Appendix C.2.

Incentives to boost generic market share Many PBMs also run mail-order pharmacies, which most often ship generic drugs. Therefore, PBMs might have additional incentives to switch patients onto generic drugs. To account for this effect we include an additional term for the total market share of generic drugs:

$$\theta_{generic} \times ms_{generic}(x_t, C_t)$$

The importance of this term depends on the mail order share of prescriptions drugs, which empirically is not very high.[62]

Carryover of inertia through doctors to new patients Finally, we modify Equation 5 to allow for a fraction κ of patients to inherit last period's choices: $x_{jt} = ms_{jt-1} \times \frac{N_{t-1}+\kappa(N_t-N_{t-1})}{N_t}$. This type of carryover could be due to doctors or social peers. A high value of κ would affect dynamic incentives for firms, especially when market size is growing.

Table A9 shows the results and compares them to the baseline analysis. We find that including a dynamic term for PBM revenue and a boost for generic market share lead to a small increase in model fit. The coefficients are also significant. However, the coefficients are also very small and do not affect other estimated coefficients. Because adding even one coefficient significantly increases computational time, we ultimately decided against including them in our main specification.

[62] According to the Kaiser Family Foundation (using 2019 data from IQVIA), patients fill about 11.6 prescriptions at retail pharmacies on average, with only 0.52 fills through mail order. Sources: https://www.kff.org/health-costs/state-indicator/number-of-mail-order-prescription-drugs-per-capita/ and https://www.kff.org/health-costs/state-indicator/retail-rx-drugs-per-capita/ (both retrieved February 17, 2023). Because mail order utilization has increased over time, the rate should be even smaller during our period of analysis.

Table A9: Parameter Estimates from Additional Specifications

		Estimates		
Parameter	Baseline	(1)	(2)	(3)
θ_s	1.70	1.71	1.72	1.81
	(0.13)	(0.26)	(0.45)	(0.26)
θ_e	1.50	1.77	1.61	1.59
	(0.05)	(0.22)	(0.35)	(0.17)
θ_j	1.11	1.24	1.13	1.29
	(0.37)	(0.17)	(0.27)	(0.17)
θ_ω	0.063	0.061	0.064	0.063
	(0.005)	(0.006)	(0.005)	(0.002)
β^{PBM}	–	0.10	–	–
	–	(0.02)	–	–
$\theta_{generic}$	–	–	0.11	–
	–	–	(0.02)	–
κ	–	–	–	0.08
	–	–	–	(0.05)
Total Error	0.458	0.435	0.439	0.456

Notes: All parameter estimates are based on the "Pct Discount" price series. Total error refers to the sum of squared errors across all branded drug-years between 1996 to 2013 (excluding periods after generic entry for each drug). Standard errors are computed using the two-step procedure outlined in the main text.

D.5 Predicted prices under 33% coinsurance counterfactual

Figure A13 presents a comparison of the equilibrium prices under the 33 percent coinsurance counterfactual to the equilibrium prices under the baseline estimated model. The coinsurance prices are higher, and exhibit greater growth over time. This is particularly obvious in the case of Lipitor, which, absent PBMs, would experience dramatic price growth in the first five years after launch.

Figure A13: Price Comparison – Baseline vs. Coinsurance

(a) Zocor

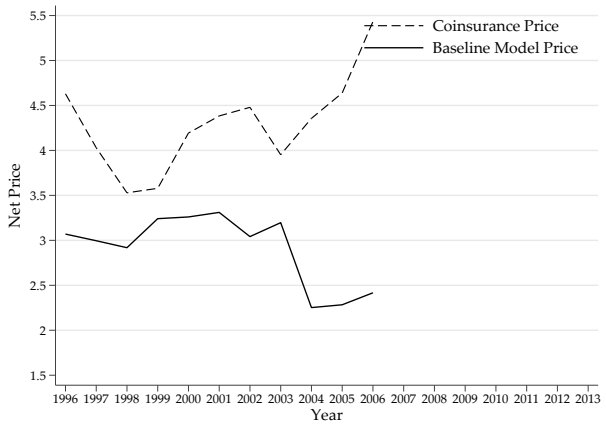

(b) Lipitor

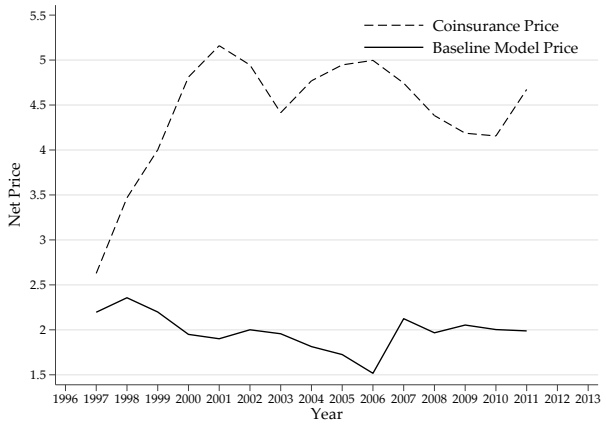

(c) Crestor

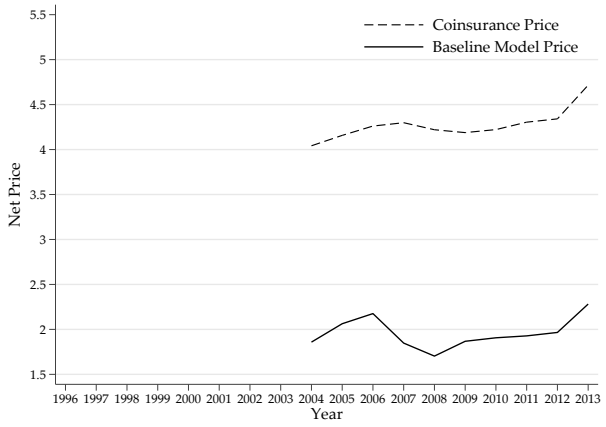

Notes: Prices predicted by the best-fit model versus the prices observed in the data.