

History-Dependent Demand and Intermediaries: Explaining Prescription Drug Pricing Dynamics

Josh Feng *

January 1, 2018

[Click here for the latest version](#)

Abstract

I show that US prescription drug prices can be better understood by accounting for history-dependent demand and industry pricing structure, which has important implications for potential government negotiation. First, I confirm the presence of inertia and history-dependent demand for generic drugs through a novel quasi-experiment. Next, I capture pricing structure in a dynamic pricing game involving drug companies and an intermediary that designs drug insurance, and estimate it using novel net price data from the anti-cholesterol market. Counterfactuals show that intermediaries reduce expenditures by over 20% through their ability to exclude drugs from insurance, but capture a significant fraction of the savings. Government negotiation would need to have a similar threat to curb the inflationary pressure from history-dependent demand.

JEL Codes: L11, L65, I12, H51

*I would like to thank my advisers Ariel Pakes, David Cutler, and Josh Lerner. In addition, I would also like to acknowledge Ernst Berndt, Lauren Cohen, Edward Glaeser, Robin Lee, Alex Mackay, Luca Maini, Anders Munk-Nielsen, and Heidi Williams for providing detailed comments and suggestions throughout the research process. Next, I would like to thank various industry contacts for providing valuable background information and data: Ryan Baum and SSR Health for providing data and suggestions on net price construction; Bob Nease and Edmund Pezalla for insights into PBM formulary decision-making. Finally, the research also benefited from comments from participants in the Harvard Industrial Organization Lunch, the Harvard Labor/Public Lunch, the Harvard Innovation and Entrepreneurship Brownbag, and the NBER Productivity Seminar.

1 Introduction

There is significant confusion surrounding the determinants of drug pricing in the United States, highlighted by Congressional and media responses to recent cases of drug inflation. These cases include drugs with no close substitutes such as Daraprim, which increased from \$13.50 per tablet to \$750 within a few weeks, and the EpiPen, which increased from \$100 to \$600 over a decade,¹ but also include cases in markets with close substitutes, such as the insulin market.² All of this has led to significant Congressional and judicial scrutiny of agents involved in the industry’s pricing structure, including drug manufacturers and pricing intermediaries, known as pharmacy benefit managers (PBMs),³ with accusations of collusion and proposed centralized price negotiation for Medicare Part D.⁴ A neglected aspect in the discussion is how prescription drug demand interacts with the pricing structure, especially given industry commentary and institutional details that point to significant history-dependence in demand patterns.

In this paper, I provide quasi-experimental evidence of history-dependent demand in several large chronic drug markets, and then highlight the role played by demand and pricing intermediaries in forming drug prices, by estimating a model of industry pricing structure on net price data and then varying demand and intermediary behavior in counterfactuals. To begin, I establish evidence of history-dependent demand by using a strategy based on discontinuities around drug launches. As a unit, patients, doctors, and pharmacists generate demand that exhibits inertia, a tendency to choose the same drug, and more general history-dependence when it comes to choosing generic or incremental drugs.⁵ To understand how demand impacts drug prices, I embed a switching-cost demand system in a model of industry pricing structure. I model the pricing structure as a finite-period dynamic game between drug companies, where they make simultaneous net price offers to a representative PBM in every period, and the PBM, in turn, sets a formulary that impacts demand for each drug. Taking the model to data on net prices in the anti-cholesterol market allows me to recover the structural parameters of the game, in particular the objectives of PBMs. Counterfactuals show that formulary structure and PBM incentives play a key role in matching the qualitative aspects of the data. They also show that drug prices and spending would be over 20% lower absent inertia or if the PBM were less averse to excluding popular drugs. These results help to explain price dynamics in one-use markets such as cures for Hepatitis C and imply that Medicare

¹See NY Times articles “Drug Goes From \$13.50 a Tablet to \$750, Overnight” and “EpiPen Price Rise Sparks Concern for Allergy Sufferers.”

²See NBC News article “Is Insulin the New EpiPen? Families Facing Sticker Shock Over 400 Percent Price Hike.”

³As I detail in Section 2, PBMs negotiate prices with drug companies, using their ability to set formularies as a tool for extracting discounts.

⁴See articles “Drug Markets Accused of Fixing Prices on Insulin” (NY Times), “Drugmakers Point Finger at Middlemen for Rising Drug Prices” (Wall Street Journal), and “Medicare Should Leverage Buying Power to Pull Down Drug Prices, White House Says” (NPR).

⁵Generic drugs contain the same active ingredient (molecule) as the corresponding branded version, and enter after the patent expires on the branded drug. Incremental drugs refer to reformulations or combinations involving the molecule in an existing drug.

Part D bargaining would need to have a significant exclusion threat to curb prices.

This paper adds to the growing literature on consumer inertia and makes three main contributions to the literature on prescription drug markets: a framework for uncovering history-dependence in drug choice, a method for constructing net prices, and a tractable model that more accurately captures industry pricing structure. There is a long literature on consumer inertia and brand loyalty, which has used random assignment, dominated choices, and discounts to identify inertia in insurance and consumer goods markets.⁶ I contribute to this literature by offering an identification strategy in settings with new products and consumer inflexibility in decision timing.⁷ In terms of prescription drug markets, my paper is the first to provide causal evidence of history-dependence in chronic drug demand, a possibly important factor in many drug company decisions.⁸ Previous studies on learning and advertising, including Crawford and Shum (2005), Dickstein (2014), Sinkinson and Starc (2015), and Lee (2016) either implicitly or explicitly find history-dependence in their demand analysis, but do so by imposing modeling assumptions. They also analyze effects over a few weeks or months, and neglect to draw out the implications for pricing.⁹ In addition, it is one of the first to use net price data in analyzing prescription drug market. Net price patterns are more reflective of market forces and I find that they can have very different dynamics to those of list prices.¹⁰ Finally, to my knowledge, my paper is the first to incorporate pricing intermediaries (PBMs) in a model of equilibrium drug pricing in US markets. Previous studies on equilibrium pricing¹¹ have typically used a direct pricing structure, but I show that PBMs play a key role in generating both quantitative and qualitative features of the net price data.

The starting point of my research is to establish causal evidence of history-dependent demand at the patient level, which I do by constructing a quasi-experiment around the launch of branded drugs. Here, history-dependence refers to past choices having a causal impact on current choice, a generalized version of inertia, which refers to people staying on the same product. Separating history-dependence from unobserved heterogeneity is important for understanding whether drug companies have incentives to invest in market share. I focus on patterns at the patient level, treating the decision as a joint one between patient, doctor, and pharmacists, although I will refer to the group as the “patient” throughout. I isolate inertia by constructing a quasi-experiment around the launches of branded drugs, which contain an ingredient (molecule) not previously available. Patients who begin treatment just after a new drug launches are able to choose it as their first choice, unlike those who begin treatment just before. Tracking the groups

⁶I discuss the literature in detail in Section 3.1.

⁷The strategy is similar in spirit to Handel (2013), who compares the insurance plan choice of new and continuing workers.

⁸See Appendix E for a detailed discussion.

⁹An exception is Ching (2010), which looks at the impact of global learning about generic quality on generic pricing and competition. Researchers have investigated pricing and inertia in other health care markets, mainly in Medicare Part D plan choice. These include Ericson (2014), Wu (2015), Fleitas (2016), and Ho et al. (2017).

¹⁰Aitken et al. (2016) collects some net price data as an additional check on their analysis.

¹¹Examples include Dunn (2012), Arcidiacono et al. (2013), and Bokhari and Fournier (2013).

over time, I find that the “after” group has a persistently higher share of patients on the new drug, even five years or more after launch. This gap means that there is a causal effect of initial choice on later choice, conditional on the assumption that patient start dates do not depend on the drugs available. This seems plausible given that doctors want to treat patients as soon as a problem arises, and I provide evidence supporting this assumption by analyzing patients who start treatment after hospitalization and by looking at smoothness in the number of new users around drug launches.

The same quasi-experimental framework permits me to show history-dependence in for generic drug demand, which is important for understanding pricing outcomes around generic entry. A key aspect of drug markets is generic entry, which occurs at the end of a branded drug’s exclusivity period.¹² Several manufacturers typically enter with generic drugs containing the same active ingredient as the branded drug, driving down its price. Factoring in generics is important for understanding market dynamics, as firms will set prices in anticipation of generic entry based on patient behavior. Using the same drug launch analysis framework, I provide both causal and descriptive evidence that patients switch at very high rates from branded drugs they are taking to generics containing the same molecule (active ingredient), but are far less likely to adopt a generic drug containing a different molecule.

I then explore the impact of these demand patterns on equilibrium pricing outcomes, by embedding them in a model reflective of the pricing structure in drug markets, most crucially the role played by PBMs. As explored in previous literature, such as Klemperer (1987) and Dubé, Hitsch and Rossi (2009), history-dependence in demand can have large impacts on pricing in an oligopoly setting, generating increasing prices over time. I build on this literature by constructing a dynamic oligopoly model that factors in history-dependent demand, while also incorporating the mediating role played by PBMs in price setting, the presence of generics, and the finite period nature of branded drugs due to patent expiration. The inclusion of PBMs in the model better captures the actual pricing structure in the industry, and is helpful in matching qualitative aspects of the price data.

To better measure the actual equilibrium pricing outcomes in drug markets, I construct novel data on net prices for anti-cholesterol drugs, finding significant differences relative to list prices. The existing literature on drug pricing generally works with data on list price, or sticker price, given the lack of data on discounts offered by drug companies. To construct data on average US net prices paid to drug companies, I use estimates from SSR Health and also collect additional revenue data from financial filings and earnings call transcripts. I divide net revenue estimates by quantity and make some minor adjustments for dosage to arrive at net prices. Unlike list prices, net prices do exhibit declines, particularly after generic entry of rival drugs. This drop is preceded by an increase in net prices before rival generic entry.

I then estimate the model using the average net price data, recovering parameters of the pricing game, in particular the PBM profit function. First, I estimate a switching-

¹²The exclusivity period ends when the patents covering the drug expire and the FDA-awarded marketing exclusivity period ends. Two major drugs have generic entry during my period of analysis.

cost model of demand on individual-level prescription drug choice and formulary data, leveraging the instrument from my earlier causal analysis to deal with unobserved heterogeneity in preferences. I then embed demand estimates in my model of pricing structure, and estimate the model using the aforementioned net price data. To do this, I play out the dynamic game involving firms and PBMs, with the former making take-it-or-leave-it offers in each period to the latter and both types of agents accounting for consumer behavior. By matching to data on net prices, I recover key parameters surrounding the profit function of PBMs. The estimates suggest that PBMs generally aim to maximize cost-effectiveness to please their customers, but are limited in their ability to exclude popular drugs and are also influenced by other factors such as ownership.

Using the model estimates, I compute counterfactual pricing dynamics under different patient and PBM behavior, with an eye towards understanding pricing in other drug markets¹³ and the potential impact of centralized Medicare Part D bargaining. Using the model estimates, I run counterfactual pricing outcomes by adjusting patient inertia and the behavior of PBMs. Equilibrium prices would be lower to start and decrease significantly over time if consumers exhibited lower levels of history-dependence or if PBM customers were less averse to formulary exclusions. This is consistent with the sharp net price decreases in drug markets with no repeat patient demand, and also suggests that Medicare Part D negotiators would need to be willing to exclude popular drugs in order to impact prices in markets with significant patient inertia.

Finally, I also compute counterfactuals under alternative pricing structures, which show that the combination of PBM incentives, discrete copay tiers, and demand inertia, are key to explaining both the quantitative and qualitative features of observed prices, and provide evidence that PBMs may curb costs in markets with demand inertia. One key feature of the net price data is that a drug with higher implied quality and a higher base of existing patients ends up being priced persistently lower relative to a drug of lower quality. A simple model of without PBM involvement is unable to capture this effect, whereas the discrete copay tier structure in my model, combined with inertia, help explain the discrepancy. The intuition is that it is hard to dislodge competitors from the formulary due to inertia and PBM incentives, and therefore drug companies need to set significantly lower prices to maintain high market shares, if they want to pursue such a strategy at all. The counterfactuals also suggest that PBMs reduce expenditures on drugs, despite their profitability.

The rest of my paper is organized as follows. Section 2 provides the relevant background surrounding drug demand, drug pricing, PBMs, the anti-cholesterol drug market, the dataset I use for my analysis, and my methodology for constructing net prices. Section 3 lays out the causal evidence on history-dependence. Section 4 introduces the dynamic pricing game, which involves drug companies and PBMs. Section 5 discusses net price calculations, model estimates, and model fit. Section 6 evaluates policy counterfactuals. Section 7 concludes.

¹³For example, the market for recently introduced Hepatitis C cures, where patients take drugs for a 12-24 week period and then never again, as they are cured. The market has seen significant reductions in net price over time.

2 Background: Drug Demand, Prescription Drug Markets, and Data

I provide an overview of key institutional details and previous findings related to prescription drug demand and prescription drug markets that motivate my research and provide context for my subsequent findings. Then, I lay out the details and market outcomes in the anti-cholesterol market, which is the focus of my structural analysis. I conclude by discussing my data sources, including my construction of net prices, and offer some motivating empirical facts based on the data.

2.1 Demand in Chronic Drug Markets - Institutional Details

To begin, I provide a brief overview of the institutional details surrounding demand in chronic drug markets. These details are conducive for my analysis of history-dependent demand and also provide some context for my findings.

2.1.1 Drug Choice Mechanics - Patients, Doctors, and Pharmacists

In this paper, I treat observed prescription drug choice as the result of a joint decision made by patient, doctor, and pharmacist. Doctors are responsible for diagnosing health problems in patients, and then make a decision on whether or not to use medication to treat the problem, as well as which drug to take.¹⁴ For each of the possible drugs, patients with insurance will face various copay amounts, which often come in two tiers for branded drugs and a separate tier for generics, rather than the full price. However, the doctor is often not aware of the details of the patient's insurance, possibly dampening price sensitivity. Forces that push back against this include insurers calling doctors to make them aware of the drugs with lowest copays on their patient's plan and doctors generally knowing that generic drugs have much lower copays.

Beyond patient and doctor, pharmacists also play a role in drug choice, by alerting doctors to insurance coverage issues and by substituting generic medication. Once a patient obtains a prescription, they usually go to the pharmacy to obtain their medication. The pharmacist they see can view the formulary on a patient's drug insurance, which usually lists copayments for various drugs. The pharmacist can choose to contact the doctor to change the prescription to a cheaper alternative, a practice known as "therapeutic substitution."¹⁵ A more common form of pharmacist behavior is to dispense generic drugs that contain the same active ingredient, or molecule, as the branded one prescribed by the doctor. This guides patients on a branded drug to switch to the generic version, which I see in my analysis. Each state in the US has generic substitution laws that govern pharmacist behavior, ranging from mandatory to optional substitution.¹⁶

¹⁴For example, in the case of a high cholesterol diagnosis, doctors can suggest changes in diet and exercise or prescribe one of many anti-cholesterol drugs.

¹⁵Therapeutic substitution without notifying the doctor is prohibited in many states.

¹⁶See the 2006 Pharmacist's Letter article "State Regulations on Generic Substitution" for details on US laws by state. Available at: <http://pharmacistsletter.therapeuticresearch.com/pl/ArticleDD.aspx?nidchk=1&cs=cepda&s=PL&pt=2&segment=1186&dd=220901>

2.1.2 Chronic Drugs and History-Dependence

In this paper, I focus on demand in chronic drug markets, which represent some of the largest drug markets and allow me to observe patient-level prescriptions over time. Chronic drug markets have the feature that once a patient starts taking medication, they keep taking medication indefinitely.¹⁷ This feature allows me to track repeated choice over time for a large number of people.¹⁸ Examples of chronic markets include treatments for high cholesterol, diabetes, multiple sclerosis, asthma, hypertension, and HIV, all of which are large markets in terms of revenue and/or number of patients both in the US and globally.¹⁹

One aspect of prescriptions in chronic drug markets is that they are valid for up to one year, creating the possibility of significant history-dependence. A written prescription from a doctor is valid for an amount of time depending on a drug's schedule. The Drug Enforcement Administration (DEA) classifies drugs into six schedules, with Schedule II drugs having the highest potential for abuse, and Schedule VI having the lowest. Painkillers such as OxyContin and attention deficit hyperactivity disorder drugs such as Adderall are examples of Schedule II drugs, while cholesterol and diabetes medication are in Schedule VI. Each state then sets its own laws on the length of prescription validity, with the shortest for Schedule II drugs and the longest for Schedule VI. For example, Massachusetts sets a 30-day supply limit on Schedule II drugs, but a one year limit on Schedule VI drugs. Therefore, patients on Schedule VI drugs can choose to refill their medication several times without having to visit their doctors.

Additional suggestive evidence of history-dependent demand comes from drug company decisions and commentary. As I detail in Appendix E, drug company decisions surrounding incremental drug entry and FDA priority review vouchers suggest that there is significant history dependence in demand. For example, they value FDA priority review vouchers, which takes four months off review time, at a significantly higher amount than what the drugs earn in those four months. In addition, drug companies regularly discuss breakdowns such as new user market share or the source of switchers to their new combination drugs.²⁰

Finally, anecdotal evidence suggests that doctors take an “if it ain't broke, don't fix it” approach to prescribing drugs, which can generate history-dependence at the patient level. From speaking with doctors and industry analysts, the conventional wisdom is that doctors are reluctant to switch patients away from an effective drug, even if the doctor thinks a different drug is superior or cheaper. This behavior can still generate patient-level history-dependence, even if the patient never switches doctor. New drugs

¹⁷Allowing for occasional gaps due to non-compliance.

¹⁸As I note later, one can also track the choices of doctors over time, but the sample is much smaller.

¹⁹“Top Therapeutic Classes by Non-Discounted Spending (US)” available on the IMS Health Top-Line Market Data page: <http://www.imshealth.com/en/about-us/news/top-line-market-data>

²⁰For example, Merck asserted in their earnings calls that Vytorin (a combination of Zocor and Zetia) was intended to attract users beyond just those already on Zocor. AstraZeneca also routinely discussed new user market share in the years after Crestor's launch. Gilead recently launched a successful HIV combination drug, Genvoya, but were pressed into acknowledging that only 8% of prescriptions were switches from people not already on some components of the combination.

can enter after a patient begins treatment or existing drugs may change copay tiers over time, changes that patients and their doctors do not respond to.

2.2 Drug Pricing - Background, Literature, and Policy

Next, I provide some background on the key issues surrounding drug pricing. This includes the literature on dynamic pricing patterns, the problems with using list price data that is common in the literature, the anecdotal evidence on the pricing structure in the industry, and the most salient policy debates and discussions.

Pharmaceutical pricing has long drawn the interest of academic researchers, who have documented dynamic pricing patterns in various drug markets, particularly near patent expiration. An early paper that presents facts on pricing in the pharmaceutical industry is Caves et al. (1991).²¹ One key finding of theirs is significant price increases in the two years leading up to patent expiration. Subsequent papers, such as Dunn (2012) on the anti-cholesterol market, have noted increasing trends in prices over time, despite competitor and generic entry. Recent work by Aitken et al. (2016) provides a broad overview of pricing trends in recent years, providing a breakdown that shows price increases on incumbent branded drugs generate a sizable fraction of the growth in drug expenditures.

A shortcoming in the literature that I address here is the use of list prices in most analyses, an especially large problem in the US market. List prices, which are known as Average Wholesale Price (AWP) or Wholesale Acquisition Cost (WAC), are akin to sticker prices for cars.²² Most payers, including public ones such as Medicaid, end up receiving an unobservable but significant discount (or rebate) off the list price or even the price paid at the pharmacy,²³ in effect paying a net price to drug companies plus some payment to intermediaries. This creates a problem when analyzing pricing dynamics, because discounts do not necessarily remain constant or move in lockstep with list prices. For example, Gilead, makers of the new generation of Hepatitis C drugs, set similar list prices in 2013 and 2014, but reported average discounts off the list price of 46% in 2014, after only offering 22% in 2013, possibly due to competitor entry and PBM exclusion threats.²⁴ A broader analysis by Bloomberg and SSR Health suggest that net prices paid to drug companies is increasing, albeit at a slower rate than list prices.²⁵ Therefore, using list prices can lead to incorrect inferences, an issue addressed here that is not accounted for in many existing analyses of prescription drug markets.²⁶

²¹A comment by Ariel Pakes on the paper advocates the use of micro panel data to understand observed pricing patterns, which is the approach taken in this paper.

²²See Berndt and Newhouse (2012) for a detailed overview of the institutional details behind these prices and the pricing system more generally.

²³Pharmacy prices are often recorded in prescription claims data.

²⁴Source: conference call transcripts available through the NASDAQ website. Also discussed in Fortune Magazine: <http://fortune.com/2015/02/04/bigger-drug-discounts-put-question-mark-over-gileads-stellar-run/>

²⁵<https://www.bloomberg.com/graphics/2016-drug-prices/>

²⁶Exceptions in the health literature include the aforementioned paper by Aitken et al. (2016).

In terms of policies surrounding drug pricing, recent controversies surrounding hyperinflation and price increases more generally have drawn the attention of policymakers, with some proposals to centralize Medicare Part D bargaining and bring more scrutiny to PBM activity. Controversial cases of hyperinflation involving Martin Shkreli's Turing Pharmaceuticals and Mylan's EpiPen have drawn Congressional scrutiny, but even areas with several competitors such as insulin markets have exhibited large price increases. As a result, some congressional members have levied collusion accusations against companies selling diabetes medication.²⁷ In addition, there have been calls for centralized Medicare Part D price bargaining, which is currently banned under federal law.²⁸ Currently, PBMs bargain on behalf of Part D plans, and they have drawn significant congressional and legal scrutiny in the past 15 years, which I detail in the following section.

2.3 Pharmacy Benefit Managers (PBMs) and the Pricing Structure in the Industry

Next, I turn to discussing PBMs, outlining their importance in prescription drug markets and the controversies surrounding them in policy debates.

The PBM industry is highly concentrated, and large PBMs are quite profitable and often owned by other entities. According to the Pharmacy Benefit Management Institute, the top three PBMs in terms of total prescription claims in 2015 were Express Scripts, CVS/Caremark, and OptumRx, totaling 73% of all prescription claims in the US.²⁹ There have been several large mergers in recent years, including mergers between Express Scripts and Medco in 2012 and OptumRx and Catamaran in 2015. They generate significant annual net revenue, as Express Scripts made \$2.5 billion in 2015.³⁰ Finally, throughout the past twenty years, both drug companies and pharmacy chains have periodically owned large PBMs. This includes the 1993 acquisition of Medco by Merck that lasted until 2003, which is relevant for my analysis of the anti-cholesterol market, as Merck owned one of the best-selling drugs, Zocor.³¹ Currently, CVS, a pharmacy chain, owns Caremark, the second largest PBM by market share.

The key role played by PBMs in the pricing system is to construct a prescription drug benefit for their customers, a key part of which is to design a formulary by assigning drugs to copay tiers or excluding them from the plan entirely. PBMs design drug plans for large health insurance companies (including Medicare Part D plans), large employers

²⁷See Washington Post article detailing the accusations by Senators Bernie Sanders and Elijah Cummings: <http://wapo.st/2fhFRb>

²⁸See blog post "How Would Government Negotiation Of Medicare Part D Drug Prices Work?" for a more detailed discussion of proposals. Available at: <http://healthaffairs.org/blog/2016/02/01/how-would-government-negotiation-of-medicare-part-d-drug-prices-work/>

²⁹See 2015 statistics from the Pharmacy Benefit Management Institute available at: https://www.pbmi.com/PBMI/Research/Industry_Research/PBMI/Research/PBMI_Industry_Research.aspx?hkey=22023612-80c4-4ada-a17e-85e7dfcbc1f8

³⁰Express Scripts garnered 26% market share in 2015, and is the clearest to analyze, as it is not part of a larger company.

³¹ See Quartz article "Big pharmacies are dismantling the industry that keeps US drug costs even sort-of under control" for a detailed overview of major PBM industry events. <https://qz.com/636823/big-pharmacies-are-dismantling-the-industry-that-keeps-us-drug-costs-even-sort-of-under-control/>

that self-insure, and other entities such as unions. Their customers usually set copay tier amounts, but the PBM is in charge of assigning drugs to different tiers. For example, a customer might set the copay tiers to \$30 and \$50, and then the PBM would assign drug A to the cheaper tier and drug B to the more expensive tier. PBMs also have the ability to severely limit access to drugs by setting stringent criteria or to exclude them from the benefit entirely, a practice that has gained traction in recent years.

PBMs profit by leveraging this benefit design power to negotiate discounts off list prices, which helps them attract more fee-paying customers. Using their power to set copay tiers for a large number of patients, PBMs negotiate with drug manufacturers to obtain discounts off the listed (gross) price. They then pass some of the savings onto customers, which makes their service more attractive. PBMs generally collect a fee per patient from their customers, but also pocket some fraction of the discounts. In addition, they profit by running mail order services, making especially high margins on generic drugs.³²

Based on anecdotal evidence, PBMs are hampered by consumer inertia both in its ability to move drugs to a higher copay tier and to exclude them, forces I allow for in my later model. In a market with several substitutable competitors, PBMs would theoretically be able to extract large savings by making drug companies bid for one favorable position.³³ However, multiple decisionmakers in the industry suggested that consumer inertia curbs their ability to extract discounts. First, moving a drug to a higher copay tier can backfire in the form of higher spending if there is significant inertia, as the same patients will still pick the drug, but the drug company may offer no discounts. Second, excluding a drug entirely may lead to complaints from PBM customers, as patients on the excluded drug are forced to switch to a different drug. This may reduce the demand for the PBM's services.

A final aspect of PBM behavior relevant to my analysis is that they are somewhat forward-looking, based on the length of their contracts and their behavior around Zocor generic entry. Based on publicly available information, PBMs sign contracts of varying lengths with customers. In particular, they sign long-term contracts with large insurers, sometimes up to 10 years.³⁴ PBM behavior around generic Zocor entry provides more direct evidence of dynamic incentives. As detailed in Aitken et al. (2009) and media coverage,³⁵ Express Scripts anticipated the launch of generic Zocor by favoring branded Zocor over Lipitor in its formulary placement, even though Zocor was significantly more expensive. Their motivation was to facilitate generic adoption, which would lead to long-run cost savings.

³²See Barron's 2005 article "Pfizer's New Headache", available at: <http://www.barrons.com/articles/SB112872627432763274>

³³This actually played out in the aforementioned Hepatitis C market.

³⁴A recent Bloomberg report highlights the 10-year contracts between Express Scripts and Anthem, as well as Catamaran and Cigna. "Express Scripts' Anthem Loss Goes Deeper Than Numbers" available at: <https://www.bloomberg.com/gadfly/articles/2017-04-25/express-scripts-anthem-loss-cuts-deep>

³⁵See Chicago Tribune article "Generic Zocor won't be a market healer" available at http://articles.chicagotribune.com/2005-12-29/business/0512290220.1_generic-zocor-lipitor-generic-version

Throughout the last two decades, PBMs have drawn significant scrutiny from policymakers for their role in increasing drug prices, in part due to the opacity of their contracts. PBM contracts with both drug companies and their customers are proprietary, making it difficult to discern the net price paid to drug companies and the share of the discounts kept by PBMs.³⁶ This has led some policymakers and drug companies to blame PBMs for increasing drug costs, including the recent controversy surrounding the price of EpiPens. For example, they claim that PBMs extract a lower net price from drug companies, but then pocket all of the discounts, which then means that insurance companies pay close to the full gross price, leading to higher premiums for patients.

In addition, regulators have repeatedly investigated PBMs for violating anti-kickback and anti-fraud statutes, with at least one lawsuit asserting foul play specifically in the anti-cholesterol market. Standard volume-based discounts are legal, but any misleading advertising by PBMs to promote drugs on which they receive the largest discounts would violate the law.³⁷ In addition, as mentioned earlier, drug companies once owned large PBMs, leading to accusations of favoritism towards their own drugs in formulary design. A prominent legal case involved the Justice Department's 2003 lawsuit against Medco for favoring drugs owned by Merck, its parent company.³⁸ The complaint, later settled, included the specific accusation that Medco favored Merck's Zocor over Pfizer's Lipitor, with the former costing more than the latter.

2.4 Features of the Anti-Cholesterol Drug Market

Here, I provide some background on the anti-cholesterol drug market, which I focus on in my analysis. The market is large in terms of revenue and number of patients, has major entry events during my analysis period, and relative homogeneity in the medical characteristics of the major drugs.

Anti-cholesterol drugs help patients manage high cholesterol, an increasingly common condition, and represent a significant share of total prescriptions and spending in the US. A recent CDC report on the market estimates that 27.9% of US adults over the age of 40 are currently taking anti-cholesterol medication.³⁹ These drugs help reduce LDL cholesterol levels, which have been linked in the medical literature to increased rates of cardiovascular disease. Given the prevalence of the condition, it is unsurprising that cholesterol drugs represent a large share of spending and prescriptions. At its peak in 2004 and 2005, cholesterol drugs represented 10% of all drug spending.⁴⁰ Even now,

³⁶Express Scripts claimed in a 2017 conference call that its clients receive 89% of rebates on average. Others argue that these statistics can be misleading, as PBMs often shroud discounts in other fees, which creates an inflated pass-through number.

³⁷See <http://www.pbmwatch.com/pbm-litigation-overview.html> for a list of major lawsuits involving PBMs and kickbacks.

³⁸See articles "U.S. Is Joining Lawsuit That Says Medco Put Profits Before Patients" in the Wall Street Journal and "Medco to Pay \$29.3 Million to Settle Complaints of Drug Switching" in the New York Times for more details. Other large PBMs such as AdvancePCS faced similar accusations at around the same time.

³⁹National Center for Health Statistics Data Brief No. 177, December 2014

⁴⁰Based on calculations from the Medical Expenditure Panel Survey (MEPS). The denominator in-

cholesterol drugs represent about 7% of all drug claims.⁴¹

There are several classes of anti-cholesterol drugs, although I will later focus on the statin class due to its market dominance and homogeneity. Prescription drugs are typically classified into classes by their biological mechanism of action, and there exist several classes within the set of anti-cholesterol drugs. The biggest sellers by number of prescriptions and revenue are Lipitor (sold by Pfizer), Zocor (sold by Merck), and Crestor (sold by AstraZeneca), which are in the statin class. Statins are typically the only drugs taken for moderate or high reductions in LDL levels, as they are more potent than alternatives in other classes.⁴² While physicians debate the differences between statins, they generally have similar efficacy and side effect profiles, with muscle pain as the most common complaint (Pedersen and Gaw 2001). My identification strategy for showing history-dependence deals with unobserved patient-drug match quality, but the scientific evidence provides additional re-assurance.

A first helpful feature of the anti-cholesterol market is the number of entry events during the period in which I have data. In Table 1, I have displayed the major drug entry events during the time period covered by the MarketScan data. Before 1996, Zocor, and to a lesser extent Pravachol and Mevacor, were the predominant cholesterol drugs on the market. The biggest brand entry events are Lipitor in 1997, Zetia in 2002, Crestor in 2003, and Vytorin in 2004.⁴³ The first generic to enter was Mevacor in 2001, followed by Zocor and Pravachol in 2006, and finally Lipitor in 2011. There are also entries of incremental drugs, such as Vytorin (combination) and Lescol XR (extended-release).

The large number of cholesterol users is another desirable feature, making the market both important and easier to study. In my MarketScan data, about 10.3 million people have a prescription for anti-cholesterol medication at some point, out of 150 million unique users. This is slightly below but roughly in line with CDC estimates of usage rate in the adult population.⁴⁴ The number of people on cholesterol medication allows me to set a very fine-grained empirical strategy for isolating history-dependence.

One remaining problematic institutional detail is that each statin comes in several dosage levels, some at different list prices, which creates comparability problems across drugs when constructing price series and constructing choice sets. Each drug is made available at different daily dosage levels, in order to reflect the desired amount of chole-

cludes all drug spending, including non-prescription drugs. The first major generic entry occurs in 2006, reducing spending.

⁴¹Similar numbers from MEPS and IMS Health Top-Line data on “Top Therapeutic Classes by Dispensed Prescriptions (U.S.)”: <http://www.imshealth.com/en/about-us/news/top-line-market-data>

⁴²Other drugs such as Zetia provide an alternative mechanisms for lowering LDL. After my data sample ends, a new class of drugs, PCSK9, has entered the market with increased ability to reduce LDL levels but surprisingly limited commercial success. This may have to do with the reluctance of patients to switch away from generic statins, a fact document in Section 3.6.

⁴³Zetia is a drug in a different class that is less potent, and Vytorin is a fixed dosage combination of Zocor and Zetia, which was popular for a brief period in the high intensity treatment class.

⁴⁴According to a 2014 CDC report, general usage rate of anti-cholesterol medication has risen from 19.9% to 27.9% between 2003 and 2012, with about a 15% usage rate in the 40-59 population. The latest census estimates in 2010 put the 45-64 age category at around 80 million people. Our data covers working adults up to age 65.

terol reduction. For example, Lipitor has 10mg, 20mg, 40mg, and 80mg options, while Zocor has 5mg, 10mg, 20mg, and 40mg dosage strengths. Drug companies often price these dosages differently. For example, Pfizer prices Lipitor 10mg at a lower price than all the other dosage levels. The major issue is that I want to compare prices on dosages that are substitutes, in order to better reflect true price relationships across drugs. A second issue arises in constructing choice sets, despite the fact that copays are always uniform across dosages within a drug and insurance plan. Some drugs may not be in a given patient's choice set at all, if it does not come in an appropriate dosage.⁴⁵

I resolve these issues by mapping drug-dosage pairs onto one of three treatment levels based on medical guidelines and generic switching patterns, and then estimate demand and construct prices for the most common level. To handle the issues listed above, I map each molecule-strength pair onto one of three target LDL-lowering levels, low/medium/high intensity, by using a 2016 report by Pharmacist's Letter/Prescriber's Letter.⁴⁶ This classification is corroborated by LDL reduction amounts reported on FDA labels⁴⁷ and by switching patterns to newly introduced generics, as most switches occur within treatment level.⁴⁸ For example, Lipitor is more potent per milligram, so Lipitor 10mg and Zocor 20mg are put in the same level. I then focus my structural analysis on the medium-intensity treatment level, which is the most common and is dominated by the three major statins (Zocor, Lipitor, and Crestor).⁴⁹

2.5 Data Sources - MarketScan, Medical Expenditure Panel Survey, and Financial Filings

Here, I provide some background on my main data sources, with a focus on the Truven MarketScan dataset, which is crucial for both my quasi-experimental analysis and for my demand estimation. I also make use of data from the Medical Expenditure Panel Survey (MEPS) for representativeness in demand and from financial filings in order to estimate average net prices.

The Truven MarketScan dataset⁵⁰ is based on medical claims data collected by Truven from large US companies. Truven offers various services to help companies manage their medical spending, and anonymizes the data for researcher use. People remain in the dataset as long as their companies still use Truven's services.

The Truven MarketScan dataset provides better patient tracking and coverage period relative to other commonly used datasets in the literature, which is crucial for my quasi-

⁴⁵For example, Zetia only comes in dosages suitable for low-intensity treatment, and Lipitor does not have a low-intensity treatment dosage.

⁴⁶<http://prescribersletter.therapeuticresearch.com/pl/ArticleDD.aspx?nidchk=1&cs=&s=PRL&pt=2&segment=3768&dd=271121>

⁴⁷Dunn (2012) displays this data in Table A1 of his online appendix

⁴⁸The idea is that generic switching right after a generic enters is based on cost considerations, and therefore patients are most likely to switch between medically equivalent dosages.

⁴⁹The lowest intensity level has many additional options beyond statins, such as Zetia and Niacin. Vytorin has significant market share in the highest intensity class from 2004-2008.

⁵⁰Details can be found at: <http://truvenhealth.com/markets/life-sciences/products/data-tools/marketscan-databases>

experimental setup. The MarketScan dataset tracks a large number of working-age users over the period from 1996-2013, ideal for setting up my empirical strategy and measuring long-term outcomes. Figure 30 shows the number of enrollees in the entire dataset, which has grown from about 2 million pre-2000 to almost 60 million in the late 2000s. In comparison, MEPS surveys overlapping waves of about 30,000 users, but only tracks each user for two years, precluding the possibility of measuring long-run patterns in drug choices. A similar dataset in nature is Medicare Part D claims data, which covers most Americans over age 65. However, many patients start on anti-cholesterol medication well before 65, so we would not have as many first-time users to analyze. In addition, Medicare Part D has only been in existence since 2006, which is after many of the major drug launches documented in Table 1.

Another key advantage of the MarketScan is the ability to infer formulary data, which is important for my demand estimation. MarketScan data comes with plan identifiers for each patient, and I infer the anti-cholesterol drug formulary for each plan from realized user choices. The size of the cholesterol market helps in that even plans covering a small number of patients have many realizations of each option, which helps rule out any selection issues based on copay and makes it easier to confidently identify formulary exclusions. I then merge the constructed formulary onto user choice data in order to estimate a demand model.

One major disadvantage for the MarketScan relative to Medicare Part D data is its lack of doctor and pharmacy identifiers, which, as I mentioned earlier, forces me to consider the patient, doctor, and pharmacist as a single decision unit. This limits my ability to delve more deeply into the importance of joint decisionmaking on demand patterns, which means I am less able to speak to policy interventions targeted specifically at patient or doctor behavior.⁵¹ Therefore, throughout the paper, patient choice implicitly refers to a joint decision by the patient and his or her doctor, although I explore how much we can learn from the data about doctor behavior in Section 3.7.

A second disadvantage is the nature of the population covered by MarketScan. As alluded to before, the dataset covers families with someone working at firms that use MarketScan to help them manage their health plans, which misses out on older and poorer individuals covered by Medicare and Medicaid, respectively, and uninsured individuals also taking anti-cholesterol medication. In addition, as shown in Table 1, the population skews male.

A final disadvantage is the significant attrition rate in users, which I address by constructing both balanced and unbalanced panels. Attrition rates from year-to-year is generally in the 20-25% range. This comes from individuals leaving firms and firms leaving the MarketScan sample. Firms come in and out of the data, based on whether they need help managing their health insurance plans, which do not suggest an obvious selection problem, but may still lead to biases in the analysis of long-term effects. I construct both unbalanced and balanced panels to look for consistency in results.⁵²

⁵¹Ellison et al. (1997) explore the role of doctors and pharmacies in a multi-stage budgeting model of demand for cephalosporins, which they use to speculate about the role of policies targeting pharmacies.

⁵²See Appendix A.1 for a detailed discussion of panel selection.

To address representativeness, I make use of claims data from MEPS in my price construction and structural analysis. MEPS surveys waves of respondents on their health-related spending and outcomes, including prescription drug claims, with each wave surveyed for a two-year period. The benefit of MEPS is that it provides weights on each respondent in order to arrive at a sample that is representative of the US market. In particular, this covers older patients not in the MarketScan dataset, and addresses other representativeness issues in MarketScan such as selection on employment. I use MEPS to construct quantity data by drug and strength, in order to arrive at more accurate market shares. I then make use of this in my structural analysis and construction of weighted list prices.

Finally, I make use of SSR Health data, company financial filings and earnings call transcripts, and IMS Health Top-Line data in order to construct estimates of average net price for each drug in a given year. As mentioned earlier, it is important to construct net prices in order to gain insight into actual market outcomes. I obtain net price estimates from SSR Health from 2007 onwards, and then replicate their analysis and extend the series backwards by collecting net revenue data from financial filings, available through SEC EDGAR, and earnings call transcripts, available through NASDAQ and Factiva. I take the revenue data and divide by sales data from MEPS in order to arrive at net price estimates.⁵³

2.6 Summary Stats - Persistence in Drug Choice

Next, I offer preliminary motivating evidence for my research from the MarketScan panels. I find a high degree of persistence in drug choice, opening up the possibility of history-dependent demand.

I find significant persistence in drug choice over time at the patient level, which of course could be driven by history-dependence or unobserved preference heterogeneity.⁵⁴ Using the MarketScan data, I track patients starting from their first anti-cholesterol drug prescription, and plot for each subsequent quarter whether they are still on the same drug, same molecule (e.g. a generic or extended release version of the drug), or a combination drug with an overlapping molecule.⁵⁵ Figure 1 summarize the patterns. 40% of people starting on branded treatments and 65% of people starting on generics are still on the same exact option after 5 years, respectively.

The data also points to more nuanced patterns in switching behavior, surrounding generics and incremental innovation. For the users starting on branded medication, they appear to switch to other versions of the same molecule, including generics and extended release. This is captured by the gap between the red and blue curve in Figure 1a. The gap between the green and the red curve captures the people who switch onto combination therapies that contain the molecule the user started on as one of its components. For

⁵³See Appendix A.2 for more details on the construction of prices and minor adjustments.

⁵⁴Unobserved demand heterogeneity can be driven by user needs, persistent doctor preferences, or even payer-related factors such as formulary design and access restrictions such as step therapy and prior authorization.

⁵⁵See Section 3 for an explicit definition of drug and molecule.

users starting on generic medication, there is negligible gap between the three curves, which suggests that, while they do switch to other options, they do not pick branded, extended release, and combination options sharing the same compound.

Finally, I also show that market shares of different drugs varies by the year in which a person starts treatment, which forms the basis of both my quasi-experimental setup and an instrument for previous choice in my demand estimation. This is demonstrated in Figure 2, where I plot market shares by starting-year cohort over time. The 2004 and 2005 cohorts started on anti-cholesterol medication after Crestor became available, whereas the 2002 and 2003 mostly started before Crestor became available. The graph shows a significant and persistent gap between the two groups, suggesting that initial conditions may matter for subsequent choice. I use this fact as the basis for both my quasi-experimental setup in Section 3 and my demand system estimation in Section 4.2.

2.7 Market Shares and Net Prices in the Anti-Cholesterol Market

In this section, I provide statistics summarizing the evolution of the anti-cholesterol market in my analysis period (1996-2013). This includes data on market shares and net prices for the medium-intensity treatment level.⁵⁶

In terms of demand, the market size grows at a large but slowing rate in the period, and market share is also significantly affected by entry events. Figure 4 summarizes market demand based on MEPS data. I compute market size in number of individuals, counting those on medium intensity dosage levels plus those diagnosed with high cholesterol but not taking medication.⁵⁷ The market size grows steadily over time, but the growth rate slows down starting around 2003.⁵⁸ The share of users on each drug generally remains pretty steady, outside of the initial entry of Lipitor and Crestor, and the impact of Zocor generics on Lipitor market share. Branded market shares drop very quickly after generics enter: Zocor drops from 20% to 4% a year either side of generic entry and Lipitor drops from 12% to a little over 1% from 2011 to 2013.

The net price data shows significantly different gross vs. net price patterns, with significant net price changes around brand and generic entry events. The inflation-adjusted gross and net price data for the three major statins is summarized in Figure 3. The list price graph, as is usual in most drug markets, shows a steady growth over time for all three drugs. However, the net price graph suggests that this is misleading, as net prices does exhibit decreases.

The three most striking features of the net price data, which my model will look to explain, are the gap in price between Zocor and Lipitor but small difference between Lipitor and Crestor, the responses in anticipation of generic Zocor entry, and the drop in prices after Zocor generics enter. Similar to the list price data, Zocor appears to have a significantly higher price relative to Lipitor and Crestor, despite being less popular

⁵⁶See earlier discussion in Section 2.4 for more detail.

⁵⁷This assumption is based on my analysis of the dosage that new patients start on, which is predominantly in the medium-intensity class from 1998 onwards.

⁵⁸A Wall Street Journal Article (“The Statin Dilemma: Merck, Pfizer Fret Over Sluggish Sales”) in 2003 discusses this slowdown.

than Lipitor. Merck does begin to offer increasingly large discounts starting in 2002, but Zocor remains significantly more expensive. Unlike the list price data, Crestor and Lipitor prices drop after Zocor generics enter midway through 2006, which may reflect a response to increased competition. Traditional models had justified list price increases after competitor drugs go generic as signs of price discrimination, where firms target the most price insensitive users. The drop in price post-generic entry follows an increase in prices on all three drugs before generic Zocor enters.

3 Causal Evidence on History-Dependent Demand - Exploiting Drug Launch Events

In this section, I test for the presence of history-dependence in the anti-cholesterol market by analyzing whether a patient starting on a drug causally impacts his or her choices in later years.

My strategy for demonstrating this causal relationship involves identifying patients who start treatment right before and after a new drug launches, and then showing persistently large differences in their subsequent drug choices. To verify the validity of the design, I provide evidence that patients in the groups are similar and also re-run the analysis in a sub-sample of users that start treatment after a hospitalization.

Using this framework, I also generate evidence on how history-dependence varies in the important cases of generics and incremental drugs, two aspects unique to pharmaceuticals. The fact that switching to a generic drug is easy for patients on the branded counterpart but hard for patients taking other drugs is particularly relevant for my structural analysis in Section 4.

I conclude with a discussion on the role doctors play in generating history-dependence. If they are the ones carrying effects across periods, then new patient choice will be influenced by previous market outcomes.

I relegate to Appendix B my results on history-dependence in other chronic drug markets and how I can use the framework to measure the long-term health impact of taking a given drug using observational data, and additional discussion of the underlying micro-foundations of the observed patterns.

3.1 Previous Literature - History-Dependence in Other Markets and Identification Strategies

Before documenting my empirical strategy, I provide a brief overview of the extensive literature on consumer inertia, and highlight similarities and differences present in my setting. Inertia typically refers to consumers staying on exactly the same product, whereas in my research, I establish inertia and additional effects on related drugs.

In prior literature, inertia has been explored in plan choice and consumer goods markets, using random initial assignment, dominated options, and discounts to gain identification. Handel (2013) explores the role of inertia in negating adverse selection in health insurance markets, and uses dominated plans and new user choice to identify

and quantify inertia. Another paper in the health insurance area is Ericson (2014), who demonstrates inertia in plan choice by taking advantage of random plan assignment for the Medicare Part D Low-Income Subsidy Program. Other papers documenting inertia in plan choice include Honka (2014) in auto insurance and Shcherbakov (2016) in cable TV plans. Outside of plan choice, Dubé et al. (2009) explore switching costs in the supermarket setting, looking at inertia after exposure to supermarket discounts. Other papers in this area include Shum (2004), who analyzes the market for breakfast cereals.

My specific setting allows me to construct an identification strategy based on a temporal discontinuity, one that may be suitable for demand in innovative markets. Temporary discounts, dominated options, and random assignment are hard to find in this setting,⁵⁹ so I construct a strategy in the spirit of new user choice used in Handel (2013). I leverage drug launches, which creates differences in product availability when patients make their first choice of a drug.

3.2 Data Setup - Quarterly Panel, Treatment and Control Groups

To begin my analysis, I construct a quarterly panel of patient drug choices and classify new patients into treatment and control groups based on when they begin treatment. My simplifying choices gloss over some details, which I discuss and account for later in robustness checks.

First, to make my analysis tractable, I convert raw prescription drug claims data into a patient by quarter dataset, recording the most frequent prescription choice of the patient in that quarter. The chosen option can be summarized by the molecule, an indicator for brand vs. generic, and an indicator for extended-release, a common form of incremental innovation.⁶⁰

This simplification of the claims to a patient by quarter dataset raises two issues that turn out to be minor: some patients use a combination of drugs and they may also switch in the middle of a quarter. Combination regimens are common in other markets such as oncology drugs, but the data suggests that they are less of an issue in the anti-cholesterol market, as only 3% of my patient-quarters contain prescription claims for different drugs in the same week.⁶¹ To assign one choice to these patient-quarters, I give each option a random ID, and take the option with the lowest ID, which should be consistent across treatment and control groups. Missing out on true switches is less of a worry, as patients usually will still be on the drug in the next quarter. In robustness tests shown in Appendix B, I deal with this issue by generating indicators from the raw data surrounding whether patients are ever prescribed a given drug, which will pick up on isolated prescriptions unaccounted for in the quarterly data.

Second, I construct a list of when patients begin treatment. For each patient, I

⁵⁹one could in theory construct a quasi-experiment based on variation in free sample or copay coupon programs

⁶⁰Extended release drugs are less popular in the anti-cholesterol market, but much more prevalent in other markets such as diabetes treatments.

⁶¹This imperfectly weeds out quarters with claims on multiple drugs where patients are actually switching.

assign them a start date based on the earliest date they have a recorded anti-cholesterol drug claim. One issue especially relevant in the unbalanced panel is that I don't see a patient's history before they enter the MarketScan panel. To mitigate the risk of mis-labeling entering users who were already using anti-cholesterol medication, I filter out users who have cholesterol claims in the first six months they are in the dataset. Mostly users have either monthly or quarterly prescription entries, but there still may be some noise introduced here. This is less of a worry in the balanced panel, as many new patients have several years of claims history.

Using this list, I sort new patients into treatment and control groups for each drug entry event. To do this, I compare each patient's start date to the list of entry event dates documented in Table 1. For my core analysis, a patient is in the control group if they start in the 180 days before a drug enters and the treatment group if they start in the 180 days after. I later test the robustness of these results by using 60-day windows, as having narrower windows helps rule out changes in environment over time, including any changes in medical guidelines and clinical evidence on drugs.⁶² The events themselves are fairly spaced out, but 40% of new patients qualify for two groups (e.g. treatment group for Zetia and control group for Crestor). For these patients, I assign them to one event randomly. Some patients also fall into no groups.

3.3 Methodology and Identification Tests

My methodology for isolating history-dependence is based on comparing patients who start medication just before and just after a new branded drug launches. It relies on the assumption that patients do not actively choose their treatment start date, which I test using smoothness and balance checks and by analyzing a subsample of patients who are presumably inflexible in their start dates. The structure provided by the methodology also allows for an alternative way to evaluate the long-term health impacts of new chronic drugs.

Using the quarterly choice data and new patient classification from Section 3.2, I track the choices of patients in the two groups to infer the causal effect of being assigned the new drug as initial treatment on the probability a patient chooses that drug in later periods. More formally, the regression specification is:

$$Y_{it} = \alpha + \beta_t Y_{i0} + \epsilon_{it} \tag{1}$$

where Y represents whether the patient chose the entering drug for the group they are in, i indexes the user, and t the quarters since launch. β_t represents the effect of initially choosing the new option on whether the patient is on the new option t quarters later.

To analyze the quasi-experiment in an intent-to-treat manner, I instrument for Y_{i0} using Z_i , an indicator for whether the patient is in the treatment group. Being in the treatment group opens up the possibility that the patient starts on the new drug, and

⁶²For example, guidelines for statin use might change in a given year, leading to a new patient cohort with different characteristics than their predecessors.

the instrumental variables framework helps recover a treatment-on-treated estimate (β_t). The framework essentially scales any gap in outcome between the groups in period t by the initial gap in outcome.

The identification assumption necessary to interpret β_t causally is that patients do not actively choose the start date of their treatment, which is reasonable given institutional details but possibly problematic. This is intuitively reasonable in a health setting, as patients probably start treatment based on medical need, on the advice of their doctors. However, there are still mechanisms that may create unobservable differences between the treatment and control groups. One possibility is that patients who end up in the treatment group anticipate future entry, and wait to start on the new drug for either poor match quality with existing treatments or financial reasons. For example, in the case of Hepatitis C, many patients with mild symptoms waited for the second entrant in the latest generation, Viekira Pak, because of the high price tag on the first entrant, Sovaldi. Another possibility is that firms often advertise heavily after launching a drug, and therefore the treatment group may be selected based on general responsiveness to advertising.

Therefore, I test this passive timing identification assumption by checking smoothness in the number of patients starting treatment, balance on observables across the two groups, and by re-running the analysis on a subset of users with concurrent inpatient events. The first two would present suggestive evidence in favor of the assumption, but still may miss out on unobservables. The third would provide more definitive evidence, as users almost certainly do not choose their inpatient admission dates based on the availability of new cholesterol medication.

For the core regressions, I set the time index t to be quarters since drug launch, but also provide robustness checks by setting it to quarters since the individual started treatment. One problem with the 180-day window described above is that patients begin treatment in different quarters. My main approach will be to compare treatment and control patients at any given point in time, indexed by quarters after drug launch, with the null hypothesis being that the two groups should make the same choices given the same environment. I also try an alternative approach where t indexes quarters since the individual started on treatment. The advantage of this approach is that it lines up all patients in the analysis sample, so that we are tracking them at the same age and experience. This change only accentuates the effects that I find, because treatment patients enter later than their control counterparts, so for any given relative period t , the new drug will be more popular for the treatment group.⁶³

3.4 Estimation Results - Inertia

I present estimates of the causal impact of starting on a new drug on the probability of choosing the same drug in later periods, what I term inertia. Assigning a patient to a new drug generates a persistently higher probability of choosing that drug in the

⁶³For example, at $t = 1$, the drug has not even launched for some control patients, whereas the new drug is already popular by the time some treatment patients enter.

future. To check the validity of the estimates, I repeat the analysis on patients who start treatment after an inpatient episode.

In Figure 5, I present two pieces of graphical evidence on inertia. The left graph in Figure 5 compares the market share on the new drug between the “before” and “after” groups, pooling together all entry events. The graph shows that there is a large initial gap between the two groups during the launch quarter, with the control group having negligible share on the new drug, and that the gap between the groups persists over time. If my aforementioned identification condition holds, then there is an equal number of patients in the control group versus the treatment group who are suitable for the new drug, but there remains a gap due to differences in initial choice conditions. Another way to view this is that the new drug becomes adopted over time at some natural rate, but switching costs are preventing suitable patients from taking the drug. The right graph shows a case study of 2007 Crestor market share by starting cohort quarter, with the vertical line representing the launch date of Crestor in late August 2003. There is a clear discontinuity in the market share, with post-launch cohorts more likely to choose Crestor, even four years later.

More formally, I present estimates of Equation (1) in Table 2, which show large effects of inertia over several years. As explained earlier, the β_t coefficient can be interpreted as a treatment-on-treated estimate of inertia. The final column in Table 2 shows that the causal effect of starting on a new branded drug on probability of subsequent usage is still 55 percent at 15 quarters after drug launch, where 100 percent would represent absolute persistence and 0 percent no inertia.⁶⁴ As mentioned earlier, previous studies of prescription drug demand had touched on dynamic effects, but none documented large effects over several years.

The coefficient estimates shown in Table 2 are not completely monotonic, which has to do with the pooled nature of the analysis combined with differential attrition in the unbalanced panel. As noted earlier, due to issues of attrition in the panel, some events have less representation in later quarters. For example, as noted above, there was high attrition in 2005. Therefore, the Crestor entry event will have differential weights in early versus late period results. I show in Table 3 that the coefficients do become monotonic once I account for attrition by using a balanced panel for my analysis.

Even accounting for attrition, the coefficients reflect a non-uniform hazard rate, an aspect I will abstract away from in my structural analysis. As shown in Table 3, β_t drops significantly in earlier periods, but then the trajectory becomes flatter in later years. This suggests that switching becomes less frequent as a patient takes it more, either because of learning or habit formation. This result is consistent with comments by PBM decisionmakers, who suggested that there is a higher rate of switching in the first six months to a year, after which switching is much less frequent. For tractability reasons, I later simplify demand to a first-order Markov process, but the actual patterns are richer.

The qualitative nature of the pooled results also holds across each entry event during the period. In Figure 19 in the Appendix, I display graphs representing the analysis for

⁶⁴See Figure 18 for a full set of β_t estimates.

major entry events in the period. The results are generally very consistent across entry events, with an initial gap between treatment and control groups that persists over time.

3.5 Establishing Design Validity - Inpatient Starters, Smoothness, and Balance

My main test of experimental validity is to re-run the same analysis on a subsample of patients who start on anti-cholesterol drugs after a hospitalization. I also show smoothness in the rate of patients starting on anti-cholesterol drugs over time, as well as balance on observables.

The key threat to validity is that some patients wait for new drugs before starting treatment, so I pick a sub-sample of patients starting treatment after hospitalization, who probably are inflexible in start timing. As described earlier, my design is no longer valid if the two groups are unobservably different in preferences, which could be the case if patients have flexibility in choosing their start time. One group unlikely to have such flexibility is hospitalized users. Those hospitalized for heart-related health events, such as heart attacks, typically start on anti-cholesterol medication right after, in order to reduce the probability of future events. Therefore, they would have little control over their start time.

I generally find similar but noisier results in this sub-sample. The sub-sample is about 6% of the full sample of patients, making the results much noisier. Table 4 reports the estimates, which are slightly smaller in magnitude than the ones estimated using the full sample, but generally similar. This bolsters the case that the estimates from the larger sample are reliable.

As additional evidence, I also look for smoothness in new users and balance on observables. Figure 6 plots the number of new patients starting on anti-cholesterol medication by month, which shows that there is little evidence of spikes near key entry events marked by the vertical lines, which start dates were sensitive to entry events.⁶⁵ In terms of balance on observables, MarketScan does not have a very rich set of characteristics, but I find that age, gender, hospitalization in the starting quarter, and number of other prescriptions are balanced across the two groups. Finally, I re-run the core analysis with age, gender, and other observable health measures as controls, and find identical estimates.

3.6 History-Dependence More Broadly - Generics and Incremental Drugs

Modifying the previous framework slightly, I provide additional evidence on how consumers behave with respect to generics and incremental drugs, two important and distinctive aspects of the pharmaceutical industry, with generics particularly relevant for my later modeling choices. Generics appear to be difficult to switch from, but easy to

⁶⁵I show the balanced panel results here because it mitigates issues surrounding user turnover in the unbalanced panel. Figure 20 in the Appendix has the corresponding plot for the unbalanced panel, and also shows no evidence of discontinuities, apart from trends within a given year, driven by users who join the dataset.

switch to for consumers on the corresponding brand drug. Switches to incremental drugs appear to be easier for users on a drug that contains a common molecule.

As mentioned in Section 2, generics and incremental drugs play a key role in many drug markets, making it important to account for them in our demand analysis, especially with an eye towards firm pricing behavior around entry of these drugs. As shown in Figure 29 in the Appendix, Zocor and Lipitor generics begin to dominate the market once they enter. In addition, Vytorin, a combination of Zocor and Zetia, enjoys a period of popularity between its launch in 2004 and the aforementioned ENHANCE trial announcement in January 2008.⁶⁶

The shortcoming of using the same framework above to analyze generic entry events is that it misses out on causal effects across types of drugs, which are also relevant for drug company behavior. The results from using the same quasi-experimental setup as above can be interpreted as the effect of starting on generic Zocor on the probability a patient chooses generic Zocor in later periods. This is alone is interesting, as a strong effect would suggest that branded drugs that enter after generic entry would struggle to convert users.⁶⁷ However, it misses out on other relevant mechanism for market analysis. This includes the impact of choosing various branded drugs on later generic adoption. For example, if Merck knows that patients on branded Zocor will switch immediately to generic Zocor, they have little incentive to build market share in the period before generic entry.

I explore additional mechanisms in two complementary ways: by breaking down the control group for generic entry events and by analyzing branded launch events that occur a little before generic entry events. The first approach is to look at the different types of choices made by patients starting just before brand and generic entry events, and how they appear to affect subsequent generic choice. For example, for the generic Zocor entry event, I can look at patients who start on Lipitor, and analyze their later adoption rates of generic Zocor. The drawback of this analysis is that the initial choice reflects unobserved heterogeneity in preferences, which will carry over into whether the patient later chooses the generic. A second, complementary approach is to analyze branded entry events that occur in the few years before the generic entry of interest, substituting in generic choice as the outcome. As discussed above, the branded entry causes treatment users to take the entering drug at a higher rate. I can then trace out the impact of starting on Crestor on later adoption of generic Zocor.

The results from the first approach show strong inertia in terms of staying on generic drugs, a high adoption rate of generics for patients who start on the same molecule, a much lower one for patients starting on other branded drugs. Panel (a) of Figure 7 shows the basic result surrounding inertia for generic Zocor.⁶⁸ Treatment users exhibit

⁶⁶For my structural analysis, I leave out Vytorin, as it is almost exclusively used for the highest treatment intensity level, and complicates both the state space and introduces the issue of multi-product firms.

⁶⁷This accentuates any price differences that already put late brand entrants at a static disadvantage, which is the focus of Gilchrist (2016). The limited commercial success of the recently launched PCSK-9 cholesterol drugs is consistent with this story.

⁶⁸Results for generic Lipitor entry are very similar.

a persistently higher market share on generic Zocor. A different way to see this is to break down the control group for branded entry events, as I do in Figure 23, which shows generic almost never switching to new branded drugs.⁶⁹

Panel (b) of Figure 7 then provides suggestive evidence of other forms of history-dependence surrounding generics, by breaking down the control group by the drug users start on before generic Zocor enters. The first noticeable feature in the graph is that over 80% of control group patients who start on branded Zocor switch over to generic Zocor within a year, suggestive of almost automatic switching. This is not surprising, given the previous discussion of generic substitution laws in Section 2.1. Another result of note is that patients who start on a different branded drug switch at much lower rates, barely reaching 20% after four years. This provides suggestive evidence that previous choice influences subsequent choice, which drug companies can use to form strategies around generic entry.

The complementary approach based on analyzing branded drug entries with generic choice as outcome confirms the existence of a causal mechanism linking brand and generic drug choices. The specific approach is to use the entries of Crestor (2003) and Vytorin (2004) to analyze generic Zocor adoption in later years. The treatment group in each case is pushed towards choosing the new entering brand. I then trace out the difference between treatment and control groups in later adoption of generic Zocor. In Figure 24, I analyze the Vytorin entry event by plotting the causal coefficient from a modified version of Equation (1), with generic Zocor choice as the outcome instead of Vytorin choice.⁷⁰ The graph shows that starting on Vytorin has a negative causal effect on the later adoption of generic Zocor, a form of history-dependence different from inertia. One interpretation of this is that Vytorin takes such patients away from branded Zocor, which later results in a more difficult switch, absent the effect of generic substitution laws.

Finally, I show using the same two approaches that incremental drugs, such as combinations, also exhibit various forms of history dependence. Here, I focus on combination drugs, including Vytorin, a combination of Zocor and a non-statin Zetia.⁷¹ First, I analyze the Vytorin entry event by breaking down the control group. The patterns in Figure 25 show that control group patients who start on a drug containing a molecule also present in the launched combination adopt at the highest rates, although still far lower than the rates observed for generics.⁷² I also take the second approach, by analyzing the impact of starting on Crestor on later usage of Vytorin. In Figure 26, I present estimates showing that there is a significant negative causal impact of starting on Crestor on later usage of Vytorin.

⁶⁹This is accentuated by the fact that for entries in this period, the only previously available generic is Mevacor, which is generally considered to be inferior to newer drugs. As mentioned above, the caveat here is there could be persistent unobserved heterogeneity in tastes or price sensitivity.

⁷⁰Similar but noisier results are present for the Crestor entry event.

⁷¹Other, much less popular combinations include Simcor (Zocor plus Niacin).

⁷²As mentioned earlier, Merck officials in earnings calls had to explicitly state that they wanted Vytorin to be adopted by all patients, not just those already on Zocor, possibly contrary to the prevailing view among investors.

3.7 The Role of Doctors

So far, I have ignored discussion of the micro-foundations driving the observed statistical patterns. Here, I focus on the role of doctors, as they can generate differences in the dynamic incentives of drug companies with respect to market share. I relegate discussion of traditional demand mechanisms to the Appendix C.3, as I will be implicitly assuming for the rest of the paper that history-dependence is a structural feature of the anti-cholesterol market.

Here, I investigate whether doctors generate additional history-dependence effects by prescribing to new patients in a way that is affected by their previous prescriptions, which would affect drug company incentives to invest in market share and hint at whether effects exist in acute markets. In my analysis far, I have focused on doctors and patients as a single decisionmaking unit. However, doctors themselves may exhibit additional inertia, in the sense that their previous prescriptions influence their current choices for new patients. In effect, doctor inertia would accentuate the impact of previous market share on current outcomes. It also matters for understanding whether similar forces may exist in acute markets, where only doctors are making repeated choices over time, each time for a new patient.⁷³

I find that doctors generally prescribe many different drugs and exhibit little variation in market shares by medical school graduation cohort, suggesting that they exhibit little additional inertia beyond the effects generated through interacting with existing patients. To gain insight into doctor behavior, I bring in publicly available Medicare Part D files, covering the 2013-2014 period. For each year, the data provides aggregate prescriptions for each doctor. As shown in Figure 9, 75% of doctors prescribe at least three different cholesterol drugs in a given year, with an average Herfindahl index of 0.41. In addition, there does not appear to be a noticeable difference in prescription behavior by medical school graduation year. Under the same conceptual framework as my empirical methodology, one might expect that doctors may stick with the most popular prescription at the time of their medical school graduation, based on how they were taught. However, there does not appear to be a difference, for example, in Crestor prescription rates for doctors graduating before and after 2003. The only noticeable difference is that recent graduates are more likely to prescribe generic medications.⁷⁴ Overall, the signs point to doctors updating current preferences based on medical guidelines, while not changing the prescriptions of existing patients.

The evidence is consistent with the story that doctors take an “if it ain’t broke, don’t fix it” to prescribing. Physicians are generally afraid of unpredictable side-effects and adverse drug interactions that emerge when patients switch, and therefore are leery of changing a patient’s prescription, even under pressure from payers.⁷⁵ Therefore, they

⁷³A leading example is the case of Hepatitis C drugs, where patients are cured after a 12-24 week drug regimen, and therefore do not make repeated choices across multiple years.

⁷⁴Anecdotally, there has been a shift to the usage generic names of drugs in medical school training to to prescribe generics when available.

⁷⁵See discussion article “Drug switching: lowering costs vs. adverse interactions, potential error” in the July/August 2000 issue of the ACP-ASIM Observer.

will only consider switching in cases where the existing drug causes side effects or if the patient has a major medical event.⁷⁶ The data does indicate that patients are more likely to switch drugs when they need a change in intensity of treatment or experience a hospitalization.

There are important caveats to note here relative to the literature on doctor prescription behavior, as the findings are not necessarily inconsistent with doctor-driven heterogeneity in prescribing patterns or effects from detailing. For my analysis, I am only interested in whether doctors causally carry previous prescription choices over to new patients, which would create a stronger gradient by medical school graduation year. Doctors can still have consistent preferences over time that leads to persistence in prescribing behavior, but this variation can average out over a cohort. Detailing can also have an effect on future prescriptions, as doctors leave returning patients on the same drug.

A final caveat is the limitations created by the data, which future research can address using more detailed data. First, I am observing behavior several years after the last big entry event, so there may have been differences across graduation cohorts in earlier years. Second, medical school graduation year may not provide a sharp instrument for studying preferences, as doctors may learn about prescribing both during medical school and later in residency. Having doctor-level micro data may be able to help identify lasting effects of initial prescribing behavior on later prescriptions, using a similar framework to the one presented here.

Given these caveats, I choose to flexibly allow for previous market shares to influence new patient choice in my structural model, which I discuss in detail in Section 4.3. The results of the structural estimation suggest that an additional doctor effect may be present, but the estimate is small in magnitude and statistically insignificant.

4 Dynamic Pricing Game - Model Setup

I now turn to understanding the interaction of demand patterns and pricing structure in determining equilibrium prices.

To do this, I first construct a finite-period dynamic game involving drug companies and PBMs, who take into account history-dependent demand in their play. In the subsequent sections, I take this model to net price data to recover parameters governing the game, in particular the PBM profit function, and then run counterfactuals varying the degree of history-dependence and PBM behavior to understand their roles in forming drug prices.

4.1 Market Structure Overview

I summarize my model of market structure in Figure 10. In the model, drug companies and a representative PBM face consumers with behavior captured by a switching-cost demand system. drug companies maximize dynamic profits by making take-it-or-leave-it

⁷⁶This story is in the spirit as the breakdown model of demand in Smallwood and Conlisk (1979).

net price offers to the PBM in each year. The PBM then pick a formulary arrangement, namely a copay tier for each drug, that maximizes its own profits. Patients then make drug choices based on the formulary they face.

In terms of characteristics of the dynamic game, I model the game is a finite-period one, which ends for each drug company once its drug goes generic. I assume perfect foresight, which helps ease computation burden, but may miss out on the impacts of unexpected changes in the environment during the period.

4.2 Demand System - Switching Costs

I capture patient behavior in a switching cost model of demand. The key aspects in the model are that consumers face copays rather than list prices and that I identify switching costs using variation in cohort means, motivated by my identification strategy in Section 3.

In the demand model, consumers are myopic, and chose either a drug or the outside option⁷⁷ in every quarter based on three factors: the quality of the options, the formulary they face, and the molecule of the drug they chose in the previous period. The first factor is represented by a quality parameter assigned to each drug and year, in order to reflect changes in medical knowledge and account for advertising campaigns. This can be thought of as the revealed-preference drug quality. Another factor is the plan formulary, which is just a set of copays associated with each option. A final factor is consumer history-dependence, which is represented by an indicator of whether the given option contains the same molecule as the option chosen by the patient in the previous period.

The assumption that only the previous period matters helps simplify the dynamic model, but loses out on some of the richness captured in the reduced form estimates. The quasi-experimental estimates shown in Section 3.4 allowed for a non-parametric inertia effect over time, and offered evidence that experience reduces switching rates. For my structural analysis, I need to simplify to a first-order Markov process for demand. Without such an assumption, I would have to keep track of market shares for each cohort, which dramatically increases the state space.

More formally, consumers have utility functions

$$u_{ijkt} = \delta_{jkt} - \alpha p_{j(f)t} + \gamma I_{j=m_{i,t-1}} + \nu I_{m_{i,t-1} \neq 0} + \xi_{ij} + \psi_{ijt} + \epsilon_{ijkt} \quad (2)$$

where i indexes the patient, j the molecule (0 is the outside option), k the form of the option (brand/generic), and t the year.

The first two terms, capturing product quality and copay sensitivity are standard in demand systems. δ is a fixed effect specific to drug, form, and year. This is allowed to fluctuate to pick up variation in advertising, changes in medical evidence, and other factors that affect the outside option such as macroeconomic conditions. p is the copay faced by the user for the given option, and implicitly depends on the formulary f , which I discuss in Section 4.4.

⁷⁷In this context, one can think about the outside option as diet and exercise.

There are two switching-related terms in the equation, one for switching between drugs and one for switching from any drug to the outside option. First, I include a binary indicator, $I_{j=m_{i,t-1}}$, for whether the choice of molecule in period $t - 1$, $m_{i,t}$, matches option j being considered in the current period. The coefficient γ will capture the switching cost by boosting the incumbent molecule, and reflect the patterns found in the quasi-experiment. Second, I include an indicator for whether the patient chose any drug in the previous period: $I_{m_{i,t-1} \neq 0}$. The coefficient ν will capture this intensive margin effect. The importance of the second term comes from the fact that while some patients do switch from a drug to the outside option, they are far less likely to pick it than a newly diagnosed patient. Including the intensive margin term helps the demand system better reflect the prevailing substitution patterns if a patient is forced to switch from their current drug, possibly due to formulary exclusions.

In terms of errors, ξ_{ij} captures any individual-specific match quality, ν_{ijt} represents any serial correlation in shocks that evolves according to $\nu_{ijt} = \rho\nu_{ij,t-1} + \zeta_{ijt}$, and ϵ_{ijkt} are idiosyncratic errors drawn from a type-I extreme value distribution (logit errors).

4.2.1 Instrumenting for Previous Choice

The main focus on this paper is on the switching cost parameter γ , and both the persistent and serially correlated shocks pose a problem in identifying the parameter. As discussed in Section 3, the interpretation we are looking for from γ is that a person choosing molecule j in the previous period will cause them to be some amount more likely to choose the same molecule in this period. This is the mechanism that needs to exist in order for firms to want to pursue dynamic strategies. The non-idiosyncratic shocks ψ and ξ will be correlated with previous choice $I_{j=m_{i,t-1}}$, and therefore a simple estimation ignoring them will lead conflate history-dependence, unobserved heterogeneity, and random persistent shocks, creating an upwards bias in γ .

I solve these issues by instrumenting using a predicted value based on the starting-cohort mean in my estimation. Given the problem outlined above, what I require is a variable that is uncorrelated with both ξ_{ij} and ψ_{ijt} , but still captures variation in the actual previous choice indicator $I_{j=m_{i,t-1}}$. My proposed solution, motivated by my quasi-experimental approach, is to the leave-one-out mean of the indicator within the same starting cohort, for each molecule j and time period t :

$$\begin{aligned}\bar{I}_{y,i,j,t} &= \frac{1}{|C_y| - 1} \sum_{k \in \{C_y \setminus i\}} I_{j=m_{k,t-1}} \\ I_{j=m_{i,t-1}} &= \tau_1 + \tau_2 \bar{I}_{y,i,j,t} + \nu_{ijt} \\ \hat{I}_{j=m_{i,t-1}} &= \hat{\tau}_1 + \hat{\tau}_2 \bar{I}_{y,i,j,t}\end{aligned}$$

where patient i is in cohort C_y if they started on treatment in year y . $\hat{I}_{j=m_{i,t-1}}$ is the value I substitute in. Since we use leave-one-out means, then this quantity will

be uncorrelated with any individual error terms, as it will only reflect errors for other individuals.

I can then use this instrument in two ways: substituting in a predicted value from the instrument for the endogenous regressor or adding the residual v_{ijt} as an additional regressor. As described in Terza et al. (2008) and outlined in work starting with Newey (1987), two common approaches for instrumenting for endogenous regressors in nonlinear models are two-stage predictor substitution (2SPS) and two-stage residual inclusion (2SRI). The 2SPS provides an analogous approach to the standard two-stage least squares routine in linear regression. The equivalent here is to substitute in $\hat{I}_{j=m_{i,t-1}}$ for $I_{j=m_{i,t-1}}$ in the logit model. The 2SRI approach, advocated for by Terza et al. (2008), involves keeping the basic specification, but adding in the residual to absorb the serially correlated error terms. I will subsequently report results using both approaches.

The requirement for the instrument to work is that there is significant variation in market shares across cohorts, a fact established in Section 3. This requirement for cohort-level variation is akin to the standard first stage relevance condition on instrumental variables. Crucially, as I showed in Figures 5b and 22, there are major difference across cohorts in choice over time, stemming from variation in drug availability and announcements of clinical trial results. Other sources of variation include direct-to-consumer advertising campaigns, which Sinkinson and Starc (2015) show have a greater effect on new users by leveraging variation in advertising supply driven by election ads.

I also take a similar approach to instrumenting for the previous period intensive margin indicator $I_{m_{i,t-1} \neq 0}$. The purpose of the term is to capture substitution patterns in the event one of the drugs becomes unavailable due to formulary exclusion. Based on raw patterns, patients already taking drugs are more likely to switch to another drug rather than the outside option. An instrument is necessary here as well, as unobservable match quality will drive patients to continually take some drug, and I use a leave-one-out mean of the rate of any drug usage in a given starting year cohort.

4.2.2 Copays vs. Prices in Demand System

Another key feature of my demand system is the use of copays, which differs from much of the previous pharmaceutical literature, and this helps add realism and yields extra insight into formulary setting agents. Recent literature analyzing prescription drug markets, including Dunn (2012), Arcidiacono et al. (2013), and Bokhari and Fournier (2013), use list price in their respective demand systems. While the simplification may capture some of the forces in the market, as copays are correlated with price, it can lead to unrealistic implications such as high marginal cost of small molecule drugs and price-sensitive consumers. Using copays leads to a more realistic structure, and opens up the possibility for understanding the forces shaping formulary design by PBMs.

One simplification in my model is the homogeneity across users, which is necessary for tractability in my dynamic game estimation. In standard demand models, parameters such as α are modeled as patient-specific coefficients, in order to reflect heterogeneity in price sensitivity. However, as I describe in Section 4.1, the state variable reflects previous period market shares. Adding heterogeneity would mean that one has to keep track of

types of individuals on each drug, multiplying the dimension of the state space. The lack of heterogeneity here is mitigated by the fact that switching costs and PBM behavior also play a large role in generating effective price elasticities, whereas in previous models, heterogeneity in α alone generated variation in price elasticities.

Finally, I use raw copay data to identify copay elasticities, and later verify that the results are comparable to those found in quasi-experimental studies. The standard concern in the literature is that using raw prices in estimating demand systems can bias the price sensitivity coefficient, as products with unobservably higher quality are priced higher. This is less of a concern here, as I identify α using variation across plans in copay differences between drugs. Although copays could reflect plan-specific unobservable differences in tastes, this is probably less likely in a prescription drug market setting, as tiers are set by PBMs covering a diverse set of customers. To check that using raw copay data yields reasonable copay elasticities, I will later compare estimates to those recovered by Einav et al. (2016), who use copay variation from the Medicare Part D donut hole to estimate copay elasticities. They generally find inelastic copay sensitivity.⁷⁸

4.3 Drug Company Pricing - Take-It-Or-Leave-It Offers and Dynamics

Next, I formalize a dynamic game between drug companies, who make net price offers every period in order to maximize dynamic profits.

In the model, drug companies submit one take-it-or-leave-it net price offers to PBMs in each period, which simplifies away from more generalized offers while still capturing the essential forces. In selecting a net price offer, drug companies take into account competitor, PBM, and patient behavior. Realized prices then represent an equilibrium outcome of the game. For tractability reasons, I choose a simpler bidding behavior, which should capture the key trade-offs, but may miss out on richer equilibria.⁷⁹

The two core components of the dynamic game are the state variable, previous market share, and the control, which is the net price submitted by drug companies to PBMs. Having previous market share as the state variable allows me to capture the history-dependent nature of demand. Drug companies trade off prices and market share today, as lower prices will lead to lower profits today but better market position tomorrow, in the form of

More formally, drug company j chooses net price offer P_{jt} in period t to maximize dynamic profits represented by the following value function:

$$V_{jt}(\mathbf{x}_t; N_t) = \max_{P_{jt}} P_{jt} N_t \sum_{k=0}^n x_{kt} Pr(m_t = j | m_{t-1} = k, f(\mathbf{P})) + \beta^{dc} V_{j,t+1}(\mathbf{x}_{t+1}(\mathbf{P}), N_{t+1})$$

⁷⁸One concern in light of the results here is that they are looking for within patient and year changes in choice, which may be diminished by the effects of inertia.

⁷⁹I discuss this in detail in Section 4.5.

\mathbf{P}_t is a vector of net price offers and \mathbf{x}_t is a vector representing the state variable, as drug companies make simultaneous offers each period and affect each other's states. I assume that each drug company owns one drug. N_t represents the market size in period t , which I assume to be exogenous to the game.

The key functions here are the mapping from price offers to formularies and the mapping from a formulary to choice probabilities. $f(\mathbf{P})$ represents the mapping from the vector of price offers to a formulary, which is based on PBM considerations. I discuss this in detail in the following section. $Pr(m_t = j | m_{t-1} = k, f(\mathbf{P}))$ represents the realized market share for company j in this period for the set of patients with previous choice k facing a formulary f . I assume the outside option is represented by $k = 0$, and there are n drugs each owned by a different drug company. β^{dc} is the drug company discount rate.

The state variable for each drug updates according to the equation:

$$x_{j,t+1} = [1 - \mu_{t+1}(1 - \kappa)]Pr(m_t = j), \forall j > 0 \quad (3)$$

so the state variable just reflects previous period market shares diluted by new patients in the subsequent period μ_{t+1} . Since I cannot pin down the role of doctors in carrying over previous market outcomes to new patients, I flexibly allow for a carryover term κ , where $\kappa = 0$ means there is no carryover, and $\kappa = 1$ means new patients behave as if they are probabilistically assigned a previous choice based on previous period market shares.

The value function contains no expectations, as I assume perfect foresight. This is a reasonable starting point, as Lipitor is almost launched at the start of the period, and Crestor is in clinical trials as of 1998. Furthermore, there are no large changes in medical evidence in the statin class. The main benefit of this assumption is that eases the computation burden, as having shocks to the system would mean computing later periods several times, once for each realization of shocks. This will miss out on effects generated by unpredictable changes in medical evidence, market size growth, and macroeconomic conditions such as the great recession.

4.3.1 Generic Drugs

One key element of the anti-cholesterol market during this period is the role of generic entry. I deal with by having generics replace their branded versions in the dynamic game once they enter, and having them be passive but still influential in the game.

One key assumption to make the model work in a non-stationary environment is that drug companies play a finite period game that ends when generics enter. Generics typically enter when a branded drug's patents expire, although some drugs are also granted FDA marketing exclusivity that could run longer. Here, I assume companies become inactive players once the patent on their drug expires at T_j , and that the terminal payoff V_{j,T_j} is a constant that I normalize to zero. I test the robustness of this assumption by also making terminal payoffs a function of market share in the last period before patent expiration.

Another key assumption I make is that existing market shares transfer from branded drugs to generics, which creates anticipatory effects. As documented in Section 3.6, over 80% of patients on a branded drug switch over to the generic once it becomes available. I make the assumption that the generic inherits the branded drug's state in its first period, and replaces the branded drug in the game. This reflects the history-dependence found in Section 3.6, and the mechanics will influence play in earlier periods, as competitors of the expiring drug will want to avoid giving the generic a large base. On the other hand, PBMs may want to move patients to the expiring branded drug to save costs in the future, as anecdotal evidence suggests Express Scripts did in anticipation of the entry of generic Zocor.

Once generics enter, I assume they only play a passive role in the pricing game. Typically, PBMs automatically place generics in a separate generic copay tier, which has a lower associated copay than the preferred tier copay. Therefore, I assume generics do not play an active role in the pricing game, and are assigned a price equal to the pharmacy price in the data, which is significantly lower than branded net prices. However, generics still affect the competitive environment for remaining branded drugs, as generics represent an option that will automatically be cheaper for patients and more attractive to PBMs.

4.4 PBM Behavior - Formulary-Setting Behavior

Finally, I model the behavior of PBMs, who maximize profits by leveraging their formulary-setting power to negotiate discounts off list price and then taking a cut of the discounts. A PBM's ability to extract large discounts is limited by its ability to move volume and the aversion of their customers to exclusions.

Given data limitations, my model simplifies the pricing structure to include a representative PBM. As alluded to earlier, the PBM industry is opaque, as contracts between drug companies and PBMs, as well as PBMs and its customers, are not publicly available. Ideally, my model would contain an oligopoly structure for the PBM industry as well, where they compete for customers and each obtain their own price offers from drug companies. This would allow for a richer set of outcomes, and allow for a full bargaining model akin to Ho et al. (2017). Instead, I include a representative PBM in my model, which one can think of as a black box representing the PBM industry. The PBM takes one price offer from each drug company, and chooses a probabilistic distribution of formularies, which looks to capture the heterogeneity in formularies observed in the MarketScan data. An equivalent interpretation is that I am modeling a continuum of non-interacting PBMs, each with some idiosyncratic profits from selecting a given formulary.

The PBM makes formulary decisions to maximize its profits, which is based on attracting customers as well as generating other sources of revenue. As mentioned in Section 2.3, PBMs profit by charging customers for their services. PBMs are typically paid a fee per patient covered, and therefore would want to maximize demand. In addition, PBMs make some margin on selling generic medication through its mail order programs, and some were owned by drug companies in the 1990s and early 2000s. This

gives them additional incentives to select a formulary that guides patients towards generics or towards drugs owned by their parent company. One challenge here is that in a more competitive environment, PBMs would not be able to favor drugs owned by their parent company. However, given the lawsuits surrounding the Merck-Medco relationship, this behavior does occur.

The key trade-off faced by the PBM in maximizing demand is to offer large discounts while still providing a generous drug benefit. In order to obtain large discounts, PBMs need to be aggressive in excluding or moving expensive drugs to higher copay tiers. However, their willingness to move drugs to higher copay tiers hinges on demand being responsive, and their willingness to exclude drugs depends on customer preferences. Their customers are typically large employers and insurance companies, who are acting on behalf of employees or members who may strongly dislike plans that exclude a drug they have been taking.

Given these forces, I now formalize the PBM profit function. Given net prices \mathbf{P} offered by pharmaceutical firms, PBMs maximize profits by choosing a probabilistic distribution over formulary arrangements $f_t \in \mathcal{F}_t$, where $\mathcal{F}_t = \{p_t, \bar{p}_t, \infty\}^{n_t}$. n_t is the number of branded drug competitors active in period t , with generic drugs automatically assigned to a lower copay tier p_g .⁸⁰ The value of p_t and \bar{p}_t are exogenous to the game.

PBMs choose the optimal formulary distribution in every period based on the following profit function:

$$K_t(\mathbf{x}_t; N_t) = \max_{f_t} N_t [W(f_t|\mathbf{x}_t) - \theta \sum_j P_{jt} s_{jt}(f_t|\mathbf{x}_t) - \chi \sum_j I(p_j = \infty) x_{jt} + M(\mathbf{s}_t(f_t|\mathbf{x}_t)) + \sigma_\omega \omega_{f_t}] + \beta^{PBM} K_{t+1}(\mathbf{x}_{t+1}(f); N_{t+1})$$

The first three terms represent demand for PBM services, which factor directly into PBM profits under the assumption that they charge a fixed fee per person they serve. As discussed informally above, PBMs want to pick a formulary that will attract the most customers.⁸¹ The first factor is welfare or consumer surplus generated by the formulary, $W(f|\mathbf{x}_t)$, which is derived from the demand system. Higher copays and exclusions will lead to lower welfare for all patients. The second term represents expected spending on anti-cholesterol drugs, with $s_{jt}(f|\mathbf{x}_t)$ representing the realized market share for drug j in period t as a function of the formulary chosen and the current state variable. Together, these two terms form a cost-effectiveness measure of the services offered by PBMs, with θ serving as the weight. The third term contains an indicator for exclusion, $I(p_j = \infty)$, and the parameter χ looks to capture any additional aversion to formularies with exclusions on the part of PBM customers.

The fourth term, $M(\mathbf{s}_t(f))$, represents a catch-all other motivations that drive PBM profits. As mentioned in Section 2.3, several large PBMs were owned by drug companies, including the Merck-Medco relationship relevant for my current study. This may create

⁸⁰This means that $|\mathcal{F}_t|$ is smaller after generics enter.

⁸¹In reality, PBMs often sign contracts with customers first, but still have reputational concerns when choosing a formulary.

an incentive for Medco to favor formularies that generate higher Zocor market share. The other highly-discussed aspect of PBMs is their profit margin on running mail-order services for generic drugs. Therefore, they may also favor formularies that lead to more generic market share. I add in flexible terms for these factors in the following way:

$$M(\mathbf{s}_t(f)) = b_1 s_{1t}(f) I(t < 2004) + b_g \sum_{j \in G_t} s_{jt}(f)$$

where b_1 represents a boost for Zocor (indexed to be $j = 1$) in the years during which it owned Medco, and b_g represents any margins PBMs make on selling generic drugs. b_g may also capture any mis-measurement in the price of generic drugs. G_t is the set of drugs that have gone generic by period t .

The fifth term, the error term ω_{ft} , captures idiosyncratic noise in the profit function of PBMs. The error term is not known to drug companies when making offers, and helps ensure the existence of an equilibrium. Without the smoothness introduced by the error, the discrete copay tier structure can generate best response functions that may never intersect, as one drug company may want to undercut another by a small amount to gain a large bump in demand. Effectively, the resulting realization will be a probabilistic distribution across potential formularies, although only small values of σ_ω can generate the pricing patterns observed in the data.⁸²

Finally, the final term captures the dynamic incentives for PBMs, who prefer certain states to others. In the model, PBMs are limited in their ability to move volume by history-dependent demand, as highlighted by the presence of x_t in the first four terms of their objective function. Therefore, some states may be more favorable for them. For example, if many patients are already on a given generic drug coming into the period, then PBMs will have an easy time generating high generic demand and pressing for higher discounts from remaining branded competitors. I flexibly allow for PBMs to have some dynamic motives through the parameter β^{PBM} .

One element missing here is the role of list prices in PBM profits, which helps simplify my model, but does not allow me to assess the common critique that PBMs prefer higher list prices. I assume in my model that PBMs are paid a fee per patient covered, but policy and industry discussion often points to PBMs receiving a cut of rebates. Rebates are the difference between list price and net price, leading to accusations that PBMs prefer very high list prices in order to inflate rebates. The challenge with including list prices is that it adds another decision to the drug company problem.⁸³ Another reason why rebates based on list price makes less sense is that PBM customers should only care about net price or net spending on drug benefits, and therefore competition between PBMs should eliminate the role of list price in payments.⁸⁴

⁸²With large σ_ω , drug companies can rely on the error term to guarantee some volume, which leads them to offer high prices.

⁸³One could also assume that list prices are exogenous, or based on demand from uninsured patients for drugs.

⁸⁴Consultants working for PBM customers argue that some employers do make the mistake of making payments to PBMs based on savings off list price.

4.5 Model Discussion - Assumptions and Missing Elements

The model I have laid out is intended to capture the interaction between components of the pricing structure in a tractable way, which means I am forced to abstract away from several important features of the industry. These include advertising in the demand analysis and the bargaining structure and competition in the PBM industry in my model of pricing structure. I discuss these factors in detail below, and aim to incorporate them in future work.

A factor not explicitly accounted for in my demand analysis is the role of advertising, an aspect of the industry that has been studied extensively. Drug companies spend a significant amount of money on advertising, in the form of detailing to doctors and direct-to-consumer advertising to patients. Previous studies on prescription drug advertising, including Anderson et al. (2012), Sinkinson and Starc (2015), Ching et al. (2017), have found that advertising plays a role in increasing market size and own market share, but can have significant market stealing effects. Other papers have documented changing promotional patterns over a drug's life cycle, including Bhattacharya and Vogt (2003) and Ellison and Ellison (2011). In my model, advertising is absorbed into the quality term, δ_{ijkt} , which shows spikes corresponding to major advertising campaigns and some drop-off for branded drugs over their respective lifecycles. However, I do not formally have advertising interact with a drug company's decision on pricing. Without this interaction, counterfactuals may overstate or understate changes in equilibrium price, depending on the strategic complementarity between pricing and advertising.

A major simplification in my model of pricing structure is the take-it-or-leave-it (TIOLI) net price offer made by drug companies to PBMs, which neglects more complicated bargaining and volume discounts. From conversations with former PBM executives, one bargaining framework used by a prominent PBM is to have drug companies submit bids for each possible formulary arrangement rather than just a uniform bid, essentially allowing drug companies to offer non-parametric volume discounts. Other PBMs jointly negotiate over sets of drugs owned by the same company in different disease areas. Finally, there may be a back-and-forth between drug companies and PBMs in some cases, which is better captured by a Nash bargaining framework, particularly the Nash-in-Nash framework in the context of oligopolies. In my model, abstracting away from volume discounts and bundled negotiations helps me simplify the pricing game played in each period. However, this can miss out on richer equilibria that can result from giving drug companies and PBMs more contracting instruments. In ongoing work, I look to add linear volume discounts, which may still miss out on bidding behavior driven by competitive considerations. In addition, the TIOLI framework is the most appropriate given my data limitations, as I do not have data on specific PBM demand that would allow me to capture the threat points in a Nash-in-Nash framework.

Another aspect not addressed directly by the model is the oligopoly structure in the PBM industry. As noted earlier, there are three large players that each cover about 25% of the market, and five that cover 90% of the market. They compete for customers, and there are significant shifts in market share over time. For example, Express Scripts started off with a small market share in the early 2000s, but grew significantly in the

subsequent years. Here, I capture competitive forces by making PBMs sensitive to customer preferences over spending and member welfare, but in reality, PBMs may compete on a variety of dimensions, including the structure of contracts and their overall formulary that includes many other drug markets. Data limitations prevent me from credibly modeling the competition in the industry in a more concrete way, which would allow for analysis of the effects of mergers and other events in the industry. I aim to address this in future research.

5 Estimation Results

In this section, I present the estimates from the model and discuss features of the model that help fit the data. I first estimate the demand system separately, recovering switching cost coefficients and revealed-preference drug qualities. Then, I plug the demand system into the dynamic game, and recover parameters of the game by matching model predictions to my net price estimates.

5.1 Demand System - Switching Costs and Copay Sensitivity

I begin by estimating the switching cost demand system separately from the rest of the dynamic game. The estimates, shown in Table 5, capture the history-dependence documented in Section 3, and highlight the importance of instrumenting for previous choice. I also find some sensitivity to copays and some variation over time in implied drug quality.

Estimates of switching costs are very large relative to copay differences, are sensitive to the inclusion of instruments and the previous intensive indicator, and are in line with the quasi-experimental estimates. The switching cost in dollar terms will be $\frac{\gamma}{\alpha}$, which captures the equivalent increase in copay that will negate the inertia. Column (2) of Table 5 display the estimates for γ without any instruments and without the previous intensive indicator $I_{m_i, t-1 \neq 0}$, which equates to about \$4.60 on a daily basis. As a comparison, gaps between copay tiers are typically less than \$1. Column 2SPS-2 uses the 2SPS approach described earlier to instrument for the two previous choice indicators, including the intensive margin term. This reduces the between-drug switching cost to about \$2.2.⁸⁵ The instrument helps absorb the additional individual-level unobservables, and the inclusion of the previous intensive indicator helps pick up on the general lack of switching to the outside option, rather than just loyalty to one particular drug. Finally, the final estimate suggests that, all else equal, the annual switching rate is around 80%, which is generally in line with the quasi-experimental coefficients.⁸⁶

The switching cost estimates are also large relative to those found in the literature, levels that would suggest large impacts on price in regular markets through invest-then-harvest type strategies. Dubé et al. (2009) find switching costs on the order of 15-19%

⁸⁵The table also contains estimates from a 2SRI approach, which yields similar results, but contains some differences in the impact of the intensive margin indicator.

⁸⁶This is just based on the CDF of the extreme value distribution, evaluated at γ .

of the purchase price, in the orange juice and margarine markets. In their simulations (under a stylized oligopoly model), they find that when switching costs reach 3-4 times the list price, average prices over a product’s lifecycle will be double relative to the case with no inertia, and will exceed the no inertia price even in the case of forward-looking consumers. This is due to a standard invest-then-harvest strategy, where firms build market share by setting low prices initially, but then raise prices to cash in on captured consumers. The nature of the market structure in my model limits this type of behavior, as I discuss in the results and counterfactuals.

The implied copay elasticities are quite small, consistent with evidence in the literature, suggesting that PBMs play a role in generating higher price elasticities. In a recent paper, Einav et al. (2016) estimate copay elasticities for various drugs, using the Medicare Part D donut hole to generate variation in copays within a patient. They find copay elasticities for Lipitor of 0.33, significantly less than 1. My estimate of α , shown in Table 5, are all between 0.6 and 0.7. Under a copay of \$1 per pill, this translates to a copay elasticity of around 0.25 for Lipitor in the years before 2007. As patients do not face prices directly, this suggests that PBMs play a role in creating a high effective price elasticity through its formulary setting behavior.

Finally, the estimation also provides drug quality estimates for each drug and year, which suggest differences in revealed preference drug quality across drugs and over time. Figure 11 provides the estimates of δ_{ijk} from Equation (2). In terms of relative quality, Lipitor has a significantly higher quality at the start of the period relative to Zocor, which lasts until generic Zocor enters. Crestor has lower implied quality than the other two. In terms of trends over time, Zocor stays relatively flat, while Lipitor and Crestor drop in quality after the generic entry of Zocor in 2006. This is due to the fact that generic Zocor steals market share rather than expanding the fraction of patients choosing a drug. Other fluctuations in implied quality may be driven by advertising campaigns, macroeconomic conditions⁸⁷, and changes in medical evidence.

5.2 Model Computation and Estimation

Before turning to the estimation results, I describe my approach to computation and estimation for the dynamic game laid out in Section 4. The inner loop of the computation plays out the pricing game under a given parameter guess, and the outer loop finds parameters that best fit the net price data, under a maximum likelihood framework.

For the inner loop, I compute the final-period value functions on a finite grid of points in the state space and then iterate back to the first period. As is standard, I construct a set of points on which to evaluate the value function. In practice, Crestor is the only active participant in the final period (2013), and therefore the only computation is a single-agent profit maximization one. Then, I take the value function from the last period in order to solve for a pricing equilibrium in the penultimate period. I proceed backwards until I reach the first period (1996).

⁸⁷As shown in Figure 4, the fraction of patients choosing the outside option (one minus the “All Drugs” number) increases during the Great Recession.

One key aspect of the model is that I allow for a continuous control, by interpolating value functions. In a given period, drug companies play a simultaneous pricing game, arriving at an equilibrium where no one wants to deviate in pricing. In practice, I start with a guess for the equilibrium price and then update the price one drug company at a time by finding their optimal price given the other prices.⁸⁸ The continuous control means that I need to be able to compute V_{t+1} under any value of \mathbf{x}_{t+1} , which I do by constructing a smooth scattered interpolant using the output from computing the value function on a discrete grid. I can then efficiently evaluate this interpolant to evaluate the value function every time I need to compute profits.

Given the inner loop structure, I then estimate the model by minimizing squared prediction error in an outer loop. For the outer loop, I search over parameters to minimize the squared error between predicted and observed prices

$$\min_{\theta, \chi, \beta^{PBM}, b_1, b_2, \kappa} \sum_t \sum_{j \notin G_t} (\hat{P}_{jt}(\theta, \chi, \beta^{PBM}, b_1, b_2, \kappa) - P_{jt})^2$$

where \hat{P}_{jt} is the model's predicted equilibrium price. This is equivalent to an assumption of uncorrelated measurement errors in the prices. The basic structure of the likelihood function does not allow for serial correlation or period-specific errors, which I look to relax in robustness checks.

I use demand estimates and additional data and assumptions to fix the other parameters in the model, including drug company discount rate, new patient fraction, market size, and average copay tiers. For demand, I plug in the estimates from column 2SPS-2 in Table 5, but also perform robustness checks using the 2SRI approach. As is standard in dynamic game estimation, I fix the drug company discount rate, in this case at $\beta = 0.89$ to reflect high user cost of capital documented in the literature.⁸⁹ Next, I use MEPS data to estimate the market size N_t and the fraction of new patients μ_t in a given year. The market size is based on the number of people taking anti-cholesterol medication or diagnosed with high cholesterol. The μ_t is based on assuming a 5% patient attrition from year-to-year, possibly from improving medical condition, and assuming that the difference $N_t - 0.95N_{t-1}$ represents the count of new patients, so $\mu_t = \frac{N_t - 0.95N_{t-1}}{N_t}$.⁹⁰ Next, I use the available MarketScan formulary data to construct $\underline{p}_t, \bar{p}_t$ in every period. I do this by taking the copay tier dollar values for each plan, and then taking the median value for the preferred and non-preferred tiers, weighting by the number of users on each plan. The use of MarketScan is necessary, as MEPS does not provide plan identifiers, and therefore I am unable to construct copay tiers. Finally, I also use median generic prices and copays from MarketScan to set $P_g = 1$ and $p_g = 0.33$.

⁸⁸I find the initial guess through a grid search, to avoid starting in a region where the company has no variation in incentives. This is present in my model because there are price ranges under which exclusion is essentially certain.

⁸⁹See estimates in DiMasi et al. (2003).

⁹⁰Attrition is required to avoid negative numbers of new patients in a given year.

5.3 Parameter Estimates and Model Fit

In this part, I provide an informal discussion of identification, discuss the results of the estimation, and aspects of the model important for fitting the observed data.

Each parameter is roughly derived from a key moment in the net price data. Although the model is non-linear, the PBM equation is simple enough to generate a rough mapping from data to parameters. θ , the weight on expected cost, is identified by the average price levels in the net price data. The less the PBM's customer cares about costs relative to welfare, the more likely drug companies will set higher prices, as they know PBMs will still put them in a favorable position. χ captures the exclusion penalty, which reflects whether higher previous market share leads to higher prices in the current period, as the threshold for exclusion is pushed up. β^{PBM} is identified based on actions around generic entry, as discussed earlier in the case of their generic Zocor strategy. b_1 is identified off any abnormal gap between Lipitor and Zocor prices in the pre-2003 period. Similarly, b_g is identified off lower brand drug price levels after generic entry, as they would face PBMs looking to switch people to generics. Finally, κ is derived from differences in pricing versus future new patient fraction μ_{t+1} . If there is no relationship, it would suggest significant carryover of previous market shares through any doctor channels.

The best fit parameter values, shown in Table 6, provide some insights into the PBMs profit function, suggesting that they generally look to design cost-effective formularies to gain customers and earn fees, but also factor in other sources of revenue. Most of the estimated parameters aim to capture PBM behavior, as we know the least about their business model and profit function. The estimation yields a estimate of $\theta = 0.973$, which is the relative weight of cost against welfare. It is statistically indistinguishable from 1, suggesting something close to a cost-effectiveness criteria. Another key parameter χ is large and significant, suggesting that PBMs have trouble excluding popular drugs. The other parameter with statistical significance is b_1 , which suggests that Medco may have had incentives to increase the market share of Zocor based on ownership.

The statistically insignificant parameter estimates suggest weaker effects from PBM dynamic incentives and not much carryover of history-dependence through doctors. The estimates show a small and insignificant value for β^{PBM} . The lack of a dynamic effect may come from the fact that PBMs may only care about this around generic entry, and do not factor in future savings on branded drugs. It is also possible that many of their contracts are short-term, and little incentive to manipulate market shares for their current customers. The carryover parameter, κ , is also small and insignificant, suggesting that drug companies dampen dynamic strategies if they know there is about to be an influx of new patients. This is consistent with the limited evidence on doctors presented in Section 3.7.

In terms of fit, the model generally captures the patterns in the net price data, including the relative prices between the drugs and the trends in prices around generic entry. Figure 12 presents a comparison of the price dynamics predicted by the model at the best-fit parameters versus the net price data presented earlier. The predictions accurately capture the absolute and relative price levels of the three main drugs, in particular the higher price of Zocor relative to Lipitor in earlier years. In addition,

it captures the trends in drug prices before and after generic entry. Prices tend to increase for all drugs before generic entry, except for Lipitor in 2006, reflecting a desire to “harvest” existing market share before competition intensifies.

The discrete copay structure in the model helps to explain the higher price of Zocor in the earlier period. As noted before, a key aspect of the net price data is that Zocor maintains a significantly higher price relative to Lipitor in the early part of the sample period. This is unusual, as Lipitor is the more popular drug and quickly leads in market share. Although having b_1 in the model helps to match the quantitative difference, I show in the counterfactuals that Zocor is still priced above Lipitor in equilibrium even with $b_1 = 0$. This is because Lipitor has a stronger dynamic incentive to gain market share, but can only do so by undercutting Zocor significantly to push it out of the preferred tier, as PBMs have difficulty satisfying existing Zocor patients. I provide an illustrative model highlighting this effect in Appendix D.

Another key aspect of the model is the PBMs ability to exclude drugs, which limits the range over which drug companies can price, in turn curbing more extreme dynamic strategies. As shown in the net price data, there are no significant price increases, contrary to the existing literature on the market impacts of consumer switching costs. This is mostly due to the threat of exclusion in the model, which limits the range over which drug companies can price, even if they have significant market share. This then has an effect on previous periods, as drug companies have less leeway to execute an “invest-then-harvest” type strategy.

As an additional check on model fit, I show that the model yields reasonable predictions for formulary outcomes and exclusion probabilities. In addition to the pricing predictions, the model also outputs the equilibrium formulary choices made by the PBM. I do not use data to discipline these predictions, as representative data on formularies is difficult to obtain, as MEPS does not include insurance plan indicators, unlike the MarketScan data.⁹¹ However, the results presented in Figure 13 are reflective of formulary trends in MarketScan and industry reports. First, there is essentially no exclusion in the early years, consistent with MarketScan data, while Crestor faces some exclusion when it first enters.⁹² Second, Zocor has a much higher probability of ending up in the high copay tier in the 1996-2006 period. A 2001 report by the California HealthCare Foundation showed 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.⁹³ Finally, after generic Zocor enters, the model predicts that PBMs will move Lipitor and Crestor to higher copay tiers and exclude them. In reality, PBMs do appear to move remaining branded drugs to higher copay tiers⁹⁴ and exclude and apply step therapy and

⁹¹MEPS does have average realized copays from patient choice, which my model does make predictions on, but the reliability of the data may present problems.

⁹²See article “US blow to AstraZeneca”: <https://www.theguardian.com/business/2003/oct/04/4>.

⁹³See Table 9 of the report “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>

⁹⁴Pfizer was forced to enter into some agreements to move back into a pre-

prior authorization to them. Although the model does not make these actions available to PBMs, they essentially make it very cumbersome for patients and doctors to obtain insurance coverage for the drugs, akin to a partial exclusion.⁹⁵

There are key limitations to the model in terms of matching pricing dynamics, which result from issues in market size computation and model specification. In chronological order, the first issue with the model is the price predictions in the 1996-1998 period. The model outputs lower quality estimates, as shown in Figure 11, as the outside option is more popular in early years. This problem may be due to an imprecise calculation of market size, as it's possible that many diagnosed patients did not even consider statin treatment.⁹⁶ A second issue surrounds the pricing level of Zocor. The data suggests a more gradual decline in Zocor net price relative to the model. The source of the discrepancy is that I assume that the Merck-Medco relationship ends entirely after 2003, when in fact it maintained some volume guarantee contracts as part of the spinoff.⁹⁷ This is solvable by adding another flexible parameter to capture the post-2003 relationship, but for my core analysis, I wanted to limit the parameter space. The third issue surrounds pricing in 2006, the year during which generic Zocor enters. In the data, all three drugs exhibit increased net prices, but my model predicts that Lipitor will set a much lower price, and raise prices in the following year. This is the result of the weaker dynamic incentives of PBMs estimated in the model, as Lipitor has more of an incentive to acquire market share than the PBM has of moving market share to Zocor. As I mentioned earlier, in reality, PBMs may only selectively pursue dynamic strategies.

6 Counterfactual Pricing Outcomes: Varying Stickiness and Intermediary Policy

In this final section, I analyze the sensitivity of market outcomes to the degree of history-dependence in demand, to changes in PBM behavior, and to elements of the market structure.

Reductions in switching costs and more PBM willingness to exclude lead to lower and declining equilibrium prices. This is consistent with net price patterns in the Hepatitis C market, and has implications for the design of potential Medicare Part D bargaining.

Removing PBMs from the pricing structure leads to significantly higher and qualitatively different prices. Combined with estimates of PBM profits, the results suggest that they do reduce overall price levels, although their impact is unclear from a social welfare perspective.

ferred tier: http://www.pfizer.com/news/press-release/press-release-detail/pfizer_and_express_scripts_sign_agreement.

⁹⁵Based on step therapy and prior authorization data from Medicare Part D formulary data.

⁹⁶Medical guidelines generally become more encouraging of statin use over this period.

⁹⁷See NY Times article "With Ties Lingering, Medco Leaves Merck" available at <http://www.nytimes.com/2003/08/20/business/with-ties-lingering-medco-leaves-merck.html>. The article details the penalties Medco would face if Zocor market share dropped below a national target.

6.1 Quantifying the Effects of History-Dependent Demand on Pricing

I begin by computing counterfactuals under alternative values of switching cost parameters, and also assess variation in other behaviors such as generic adoption.

To begin, I evaluate the equilibrium price dynamics under a scenario with no drug-specific switching costs, and find price deflation. To implement this, I run the game under parameters (γ', ν') , setting $\gamma' = 0$ and transferring its effect over to $\nu' = \nu + \gamma$. This helps preserve the intensive margin versus extensive margin, but removes a drug-specific demand advantage. I also recompute implied drug qualities in each period as an additional way to maintain market shares. Figure 14 presents a comparison of prices under the counterfactual versus those from the baseline model. Prices are generally much lower in level, but also exhibit deflation in the period before 2006. This is due to the increasing levels of competition, as each drug can no longer rely on an existing patient base. The prices in later years are only slightly reduced, as lower inertia also hurts the competitiveness of generic Zocor. All of these effects would be accentuated if I also simultaneously set $\chi = 0$, equivalent to having the PBM customers ignore drug-specific exclusions. In the results shown, I still keep the exclusion effect, which still provides a boost to equilibrium prices. Overall, the results suggest that inertia plays a significant role in preventing deflation in drug prices when competitors or generics enter.

Another aspect of history-dependent demand is generic adoption, but counterfactuals show this to have limited impact on pricing strategies. To do this, I modify Equation (3), which governs state updating, to allow for some fraction h of projected share to be retained by the generic, with $h = 1$ nesting the current case and $h = 0$ nesting the case where generics enter without any incumbent advantages:

$$\tilde{x}_{j,t+1} = h[1 - \mu_{t+1}(1 - \kappa)]Pr(m_t = j), j \notin G_t, j \in G_{t+1}$$

where G_t once again represents the set of drugs that have gone generic by period t . Therefore, this update modification only applies across periods when there is a transition from brand to generic. Figure 31 in the Appendix presents the predicted prices in the case $h = 0$, when patients do not exhibit loyalty to molecule. The predicted prices are almost identical, which suggests that generics are inherently competitive to begin with, so making them more competitive does not affect branded drugs by much, as the companies owning them focus on preserving their current patient base. The caveat here is that, for internal consistency, one would need to also consider the role of the expiring branded drug if it does not automatically lose most of its patients.

6.2 Changes to PBM Behavior

Next, I evaluate counterfactuals under different PBM behavior. This yields some insights into the necessary parameters of potential Medicare Part D price negotiation that would yield desired pricing outcomes.

First, I compute a counterfactual with no PBM customer aversion to exclusion, $\chi = 0$, which results in lower price levels and price deflation, suggesting that potential Medicare Part D bargaining would need to have a strong exclusion threat to curb prices. Figure 15

presents the results under this alternate scenario. The equilibrium price levels are consistently lower than those predicted by the baseline model, apart from the initial period pricing of Lipitor and Crestor. This reflects the importance of existing market share for bolstering a drug company’s negotiating position relative to PBMs. Prices also exhibit a downward trend as more competition enters, with Lipitor prices dropping towards the generic Zocor price of \$1. With a stronger exclusion threat, drug companies offer more competitive bids, and therefore competitive forces dominate in this counterfactual. This result suggests that Medicare Part D would need to have a strong exclusion threat for popular drugs in order to curb prices. This runs counter to existing Medicare Part D rules, which have mandates for inclusion of drugs, and may also be politically infeasible.

Next, I evaluate counterfactuals under which there is no Merck-Medco relationship, finding a smaller but still present Zocor-Lipitor price differential. As discussed in Section 5.3, the presence of the b_1 parameter is necessary to capture the sizable difference between Zocor and Lipitor prices. However, as shown in Figure 16, the qualitative effects remain after setting $b_1 = 0$. This reflects the intuition discussed in Appendix D that Lipitor needs to undercut Zocor in order to gain and maintain its large market share, which helps it weather future competition.

There are potentially many other interesting counterfactuals here, based on the components of the model. This includes PBMs intentionally restricting the choice set \mathcal{F} , for example by committing to only having one drug in the preferred tier. The model also allows for the flexible addition of other profit motives, or variation in PBM behavior over time.⁹⁸ This includes varying the amount PBMs make on running their mail order services for generic drugs.

6.3 Do PBMs Reduce Drug Expenditures?

Finally, I look to analyze the impacts of changes in the market structure on pricing outcomes. My main counterfactual will be to look at the overall impact of PBMs on drug expenditures. I also run additional counterfactuals to highlight the importance of exclusion relative to copay tiers in curbing expenditures.

For my primary counterfactual here, I remove the PBM from the model and have patients face a fixed fraction of prices set by drug companies. This structure is akin to a coinsurance structure, where insurance companies still provide some insurance, but PBMs are not involved to provide an exclusion threat.⁹⁹ For the exercise, I fix a coinsurance rate $r = 0.33$, which is larger than the observed net price to copay ratio in the data and even typical coinsurance rates. A smaller ratio would lead to higher equilibrium prices, and lead to a larger estimate of PBM impact.

The results from the exercise show much larger swings in prices and a generally higher price level. Figure 17 shows the counterfactual price dynamics under the modified model. In general, price levels are significantly higher, as there is no longer an exclusion

⁹⁸Unexpected changes would require a more complex computation, but as there is no estimation, the computation should be manageable.

⁹⁹The motivation for moving from a copay structure to a coinsurance one is that without the threat of exclusion, drug companies will price at a very high level and be happy to stay in the high copay tier.

threat from PBMs. In addition, the prices exhibit larger swings, reflective of the larger pricing range available to drug companies. For example, Lipitor sets an extremely low entry price to capture market share initially, and then raises prices from under \$1 to \$4 within two years. In a copay tier structure, this strategy would be less appealing, as low prices may not yield much additional demand as the drug remains in the low copay tier, and high prices would lead to exclusion rather than profitable “harvesting” of market share.¹⁰⁰ The model also predicts bumps in prices before the Lipitor, Crestor, generic Zocor, and generic Lipitor entry events and subsequent declines after. This also captures some degree of “harvesting” by all companies.

More specifically, the change in structure alters the relative prices of the drugs, particularly Lipitor and Zocor. As mentioned earlier, the sustained higher price of Zocor relative to Lipitor is one of the most interesting features of equilibrium prices in the market, and, as discussed in Section 5.3, the structure in my model helps to explain this qualitative result. Without the structure, Lipitor is priced more in accordance with its dominant market position and edge in static demand.

The results of the counterfactual also provide an estimate on whether PBMs increase or decrease overall drug expenditures. Table 7 shows that net revenue earned by drug companies over the 18 year period would be \$195 billion under the structure without PBMs, as opposed to \$151 billion in the data. This number does not include payments made to PBMs, which is the other relevant component in assessing PBM impact on drug expenditures. As noted earlier, Express Scripts made \$2.5 billion in 2015, with a 25% market share. This equates to \$18 billion in PBM revenue over 18 years, projecting Express Scripts out to the whole market and assuming that 10% of the profits are from the anti-cholesterol market.¹⁰¹ This represents something of an upper-bound, as PBMs have become more profitable over time, and the anti-cholesterol market has stayed at or below 10% of all drug expenditures for the period. Even at the upper-bound, PBMs still reduce total expenditures on anti-cholesterol drugs by about 15%.

Two caveats are important to note here: I am not making any welfare claims about the impact of PBMs and I am also not ruling out that some direct pricing model cannot generate the patterns in the net price data. In terms of welfare, a key aspect of drug pricing debates is the role of profits in encouraging investments in new therapies. As such, paying \$44 billion more to drug companies instead of \$18 billion to PBMs, as is the differences between the structures analyzed here, may lead to longer-term benefits. Second, the fact that my model does not allow for heterogeneity in demand opens the door for a richer model without PBMs to explain prices. For example, it could be that Zocor attracts price-insensitive consumers, and therefore exists as a high-end product relative to mass-market Lipitor. This is an implausible explanation, as characteristics such as income, age, and gender are very similar across the group of patients on each drug.¹⁰²

¹⁰⁰From the standard invest-then-harvest strategy discussed earlier.

¹⁰¹Other ways to estimate this, including using a per prescription profit from PBM financial filings also leads to similar estimates.

¹⁰²Based on MEPS data.

Finally, I also assess additional counterfactuals related to direct pricing and formulary structure. These results are summarized in Table 7. First, I assess demand under a case where patients are uninsured and faced prices directly, under the assumption that their demand parameters remain the same. This unsurprisingly leads to much lower demand and correspondingly lower prices on the three drugs, as only patients with the greatest need will use the drugs. Second, I play with variants of the copay structure and formulary choice set. Having a multi-tiered formulary produces \$8 billion savings, but the effects are small, given the insensitivity of patients to copays. PBMs committing to having only one preferred drug also reduces expected spending, but the savings are much smaller relative to a case where they commit to only including one branded drug.

7 Conclusion

In this paper, I have provided two frameworks that help us better understand prescription drug markets and assess potential drug pricing policy. The results here show that there are significant history-dependence in drug demand, that modeling the pricing structure is the key to understanding market outcomes, and history-dependence and PBM incentives can have a large impact on equilibrium price dynamics.

My first contribution is to provide a quasi-experimental framework for confirming that there are large effects from inertia and other forms of history-dependence in chronic drug demand, which is important to account for in analyses of the industry and suggests future work into the micro-foundations. My estimates suggest that inertia has significant effects on choice probabilities over several years, which I then use to analyze pricing outcomes. However, as I detail in Appendix E, there are many firm investment and entry decisions that also appear to be affected by history-dependence in demand, suggesting that most research into the industry should take these effects into consideration. It also points to future research into the effect of demand on innovation, which can help guide market and innovation policies related to the pharmaceutical industry. In addition, my research offers a guess as to the micro-foundations of the observed patterns, namely doctors that take an “if it ain’t broke, don’t fix it” approach to chronic drug prescription, but this can be explored using richer data.

The quasi-experimental framework provides insight into dynamic linkages in choices by leveraging a temporal discontinuity, an approach that can be applied in other pharmaceutical settings and other industries. Here, I focus on variation generated by drug launches to alter initial conditions, under the assumption that patients do not actively choose the time they begin treatment, a more likely and verifiable assumption in prescription drugs. As discussed in the context of the ENHANCE trial, other significant events can also be analyzed using this framework. Beyond pharmaceuticals, it is a little harder to find settings where start timing is exogenous, but institutional details that limit choices by age serve as a prime candidate for analyzing other dynamic linkages in choices and outcomes.

In the context of prescription drugs, my framework can be used to analyze the health and spending impacts of new drugs. As briefly explored in the Appendix C.2, the

framework can be used to assess the long-term health and spending impacts of a new drug relative to the existing standard of care. In the case of anti-cholesterol drugs, this analysis yields little health differences, consistent with clinical trial evidence. However, the methodology may help in other contexts in using observation health data to evaluate causal treatment effects for drugs.

A more minor contribution here is the novel construction and usage of net price estimates for understanding the forces in prescription markets, which can be extended for a broader analysis of the industry. In the paper, I construct net prices using net revenue data reported in financial filings and earnings call transcripts. This data has become more readily available over time, and drug companies generally report net revenues by country or region for their best selling drugs in recent years. Net prices differ significantly from list prices, and likely capture the forces in the market more accurately, as list prices tend to just increase steadily over time. Future research should make use of net prices in order to better capture actual outcomes in the industry.

The second major contribution of the paper is a more concrete model of pricing structure in prescription drug markets, which helps to explain equilibrium price outcomes and offers a basis for future research. My model is the first to include a pricing intermediary, the pharmacy benefit manager, in an analysis of the pharmaceutical industry. PBMs take price offers from drug companies, and set a formulary that is faced by patients in accordance to their profit motives. The discrete nature of copay tiers and the PBM's profit motives help the model explain the salient features of equilibrium prices in the industry. As shown in the direct pricing counterfactual, a change in market structure creates major quantitative and qualitative differences in pricing outcomes, which may be difficult to explain even with a more complex direct pricing model.

In terms of pricing policy and implications, the results of the estimation suggest that lower inertia will lead to significantly deflation in prices as competitors enter, Medicare Part D bargaining would need to have a credible exclusion threat to curb prices, and that PBMs, on net, reduced drug expenditures in the anti-cholesterol market. The counterfactuals presented in Section 6 simulate outcomes under different patient and intermediary behavior, as well as variation in pricing structure. Inertia plays a significant role in preventing deflation as competing drugs and generics enter, so any changes in patient behavior may lead to significant effects on prices. The counterfactuals varying PBM behavior suggest that if Medicare Part D replaced them with a central bargaining framework, the decisionmaker would need to be able to exclude popular drugs in order to curb net price growth, something that runs counter to current inclusion rules and possible political constraints. Finally, the framework provides a simplified way to assess the impact of PBMs on drug expenditures, which suggests that despite their profitability, they do generate significant reductions in net drug expenditures.

Finally, I aim to use the framework to formally analyze investment behavior and innovation policy in the industry. Various policies exist to encourage investment in drug development, and some of them, such as rare disease vouchers, exclusivity on combination drugs, and incentives to develop antibiotics, may depend on features of demand and pricing structure. In ongoing work, I look to use the framework as a basis

for analyzing the investment behavior of drug companies with respect to incremental drugs. The framework needs to be extended to allow for multi-drug firms, but the mechanics are built in to analyze the impact of history-dependence in incremental drug adoption on the value to firms from introducing incremental drugs. This value may be especially sensitive to entry timing, as entering right before generic or competitor entry would allow enough time to build market share needed to maintain prices in the face of competition.

References

- Aitken, Murray, Ernst R. Berndt, and David M. Cutler**, “Prescription drug spending trends in the united states: Looking beyond the turning point,” *Health Affairs*, 2009, 28 (1).
- , — , **David Cutler, Michael Kleinrock, and Luca Maini**, “Has the era of slow growth for prescription drug spending ended?,” *Health Affairs*, 2016, 35 (9), 1595–1603.
- Anderson, Simon, Federico Ciliberto, Jura Liaukonyte, and Regis Renault**, “Push-me pull-you: comparative advertising in the OTC analgesics industry,” 2012, (43383).
- Arcidiacono, Peter, Paul B. Ellickson, Peter Landry, and David B. Ridley**, “Pharmaceutical followers,” *International Journal of Industrial Organization*, 2013, 31 (5), 538–553.
- Berndt, Ernst R. and Joseph P. Newhouse**, *Pricing and Reimbursement in US Pharmaceutical Markets* number April 2017 2012.
- Bhattacharya, Jayanta and William B. Vogt**, “A Simple Model of Pharmaceutical Price Dynamics,” *Journal of Law and Economics*, 2003, XLVI, 599–626.
- Bokhari, Farasat A S and Gary M. Fournier**, “Entry in the ADHD drugs market: Welfare impact of generics and me-too’s,” *Journal of Industrial Economics*, 2013, 61 (2), 339–392.
- Caves, Richard E., Michael D. Whinston, Mark a. Hurwitz, Ariel Pakes, and Peter Temin**, “Patent Expiration, Entry, and Competition in the U.S. Pharmaceutical Industry,” *Brookings Papers on Economic Activity. Microeconomics*, 1991, 1991 (1991), 1.
- Ching, Andrew T**, “A Dynamic Structural Oligopoly Model for the Prescription Drug Market After Patent Expiration,” *International Economic Review*, 2010, 51 (4), 1175–1207.
- Ching, Andrew T., Robert Clark, Ignatius Horstmann, and Hyunwoo Lim**, “The Effects of Publicity on Demand: The Case of Anti-Cholesterol Drugs,” *Marketing Science*, 2017, p. 150817104945009.
- Crawford, Gregory S and Matthew Shum**, “Uncertainty and Learning in Pharmaceutical Demand,” *Econometrica*, 2005, 73 (4), 1137–1173.
- Dickstein, Michael**, “Efficient provision of experience goods: Evidence from antidepressant choice,” 2014.

- DiMasi, Joseph A., Ronald W. Hansen, and Henry G. Grabowski**, “The price of innovation: New estimates of drug development costs,” *Journal of Health Economics*, 2003, *22* (2), 151–185.
- Dubé, Jean-Pierre, Günter J Hitsch, and Peter E Rossi**, “Do Switching Costs Make Markets Less Competitive?,” *Journal of Marketing Research*, 2009, *46* (4), 435–445.
- Dunn, Abe**, “Drug Innovations and Welfare Measures Computed from Market Demand: The Case of Anti-Cholesterol Drugs,” *American Economic Journal: Applied Economics*, 2012, *4* (3), 167–189.
- Einav, Liran, Amy Finkelstein, and Maria Polyakova**, “Private Provision of Social Insurance: Drug-Specific Price Elasticities and Cost Sharing in Medicare Part D,” 2016.
- Ellison, Glenn and Sara Fisher Ellison**, “Strategic entry deterrence and the behavior of pharmaceutical incumbents prior to patent expiration,” *American Economic Journal: Microeconomics*, 2011, *3* (1), 1–36.
- Ellison, Sara Fisher, Iain Cockburn, Zvi Griliches, and Jerry Hausman**, “Characteristics of Demand for Pharmaceutical Products: An Examination of Four Cephalosporins,” *Source: The RAND Journal of Economics RAND Journal of Economics*, 1997, *28* (3), 426–446.
- Ericson, Keith M.**, “Consumer inertia and firm pricing in the medicare part d prescription drug insurance exchange,” *American Economic Journal: Economic Policy*, 2014, *6* (1 B), 38–64.
- Fleitas, Sebastian**, “Dynamic Competition and Price Regulation When Consumers Have Inertia : Evidence from Medicare Part D,” 2016, pp. 1–70.
- Gilchrist, Duncan S**, “Patents as a Spur to Subsequent Innovation? Evidence from Pharmaceuticals,” *American Economic Journal: Applied Economics*, 2016, *8* (4), 189–221.
- Handel, Benjamin R.**, “Adverse Selection and Inertia in Health Insurance Markets :,” *American Economic Review*, 2013, *103* (7), 2643–2682.
- Ho, Kate, Joseph Hogan, and Fiona Scott Morton**, “The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program,” 2017.
- Honka, Elisabeth**, “Quantifying search and switching costs in the US auto insurance industry,” *RAND Journal of Economics*, 2014, *45* (4), 847–884.
- Huckfeldt, Peter and Christopher Knittel**, “Pharmaceutical Use Following Generic Entry: Paying Less and Buying Less,” *NBER Working Paper*, 2011.

- Klemperer, Paul**, “The Competitiveness of Markets with Switching Costs,” *RAND Journal of Economics*, 1987, 18 (1), 138–150.
- Lee, Chung-Ying**, “Brand Loyalty and Learning in Pharmaceutical Demand,” *Academia Economic Papers*, 2016, 44 (3), 333–359.
- Newey, Whitney K**, “Efficient Estimation of Limited Dependent Variable Models with Endogenous Explanatory Variables,” *Journal of Econometrics*, 1987, 36, 231–250.
- Pedersen, Terje and Allan Gaw**, “Statins—similarities and differences.,” *The American journal of managed care*, 2001, 7 (5 Suppl), S132–7.
- Ridley, David B. and Stephane A. Régnier**, “The commercial market for priority review vouchers,” *Health Affairs*, 2016, 35 (5), 776–783.
- Shcherbakov, Oleksandr**, “Measuring consumer switching costs in the wireless industry,” *RAND Journal of Economics*, 2016, 47, 366–393.
- Shum, Matthew**, “Does Advertising Overcome Brand Loyalty ? Evidence from the Breakfast-Cereals Market,” *Journal of Economics & Management Strategy*, 2004, 13 (2), 241–272.
- Sinkinson, Michael and Amanda Starc**, “Ask Your Doctor? Direct-to-Consumer Advertising of Pharmaceuticals,” *NBER Working Paper 21045*, 2015.
- Smallwood, Dennis E. and John Conlisk**, “Product Quality in Markets Where Consumers are Imperfectly Informed,” *The Quarterly Journal of Economics*, 1979, 93 (1), 1–23.
- Terza, Joseph V., Anirban Basu, and Paul J. Rathouz**, “Two-stage residual inclusion estimation: Addressing endogeneity in health econometric modeling,” *Journal of Health Economics*, 2008, 27 (3), 531–543.
- Wu, Yufei**, “Supply Response to Consumer Inertia: Strategic Pricing in Medicare Part D,” *Job Market Paper*, 2015, pp. 1–49.

8 Tables and Figures

Drug	Approval Date	First Date in Claims	Drug Company
Mevacor	8/31/1987	1/1/1996	Merck
Zocor	12/23/1991	1/1/1996	Merck
Pravachol	10/31/1991	1/1/1996	Bristol Myers Squibb
Lescol	12/31/1993	1/1/1996	Novartis
Lipitor	12/17/1996	1/31/1997	Pfizer
Niacin	7/28/1997	9/25/1997	Abbvie
Baycol*	6/26/1997	1/29/1998	Bayer
Lescol (XR)	10/6/2000	12/26/2000	Novartis
Mevacor (G)	12/17/2001	12/20/2001	-
Advicor [Lovastatin/Niacin]	12/17/2001	1/24/2002	Kos Pharma
Mevacor (XR)	6/26/2002	7/30/2002	Covis
Zetia	10/25/2002	11/6/2002	Merck
Crestor	8/12/2003	8/20/2003	AstraZeneca
Vytorin [Zocor/Zetia]	7/23/2004	7/29/2004	Merck
Pravachol (G)	4/24/2006	4/24/2006	-
Zocor (G)	6/23/2006	6/23/2006	-
Simcor [Zocor/Niacin]	2/15/2008	3/12/2008	Abbvie
Livalo	8/3/2009	6/28/2010	Kowa/Eli Lilly
Lipitor (G)	11/30/2011	11/30/2011	-
Lescol (G)	4/11/2012	4/16/2012	-
Liptruzet [Lipitor/Zetia]	5/3/2013	5/8/2013	Merck

Table 1: Summary of entry events during the 1996-2013 period, with the first date in which the drug appears in MarketScan data and the drug company that owns the drug also reported. There is often a small delay between approval and subsequent availability to patients, due to logistical issues. The major events in the sample period (1996-2013) involve Zocor, Lipitor, and Crestor, and are highlighted in bold. (G) refers to generic and (XR) refers to extended release. Drugs with bracketed information are fixed-dosage combination drugs. Baycol was pulled from the market in 2001 due to safety issues. Sources: Drugs@FDA and drugs.com

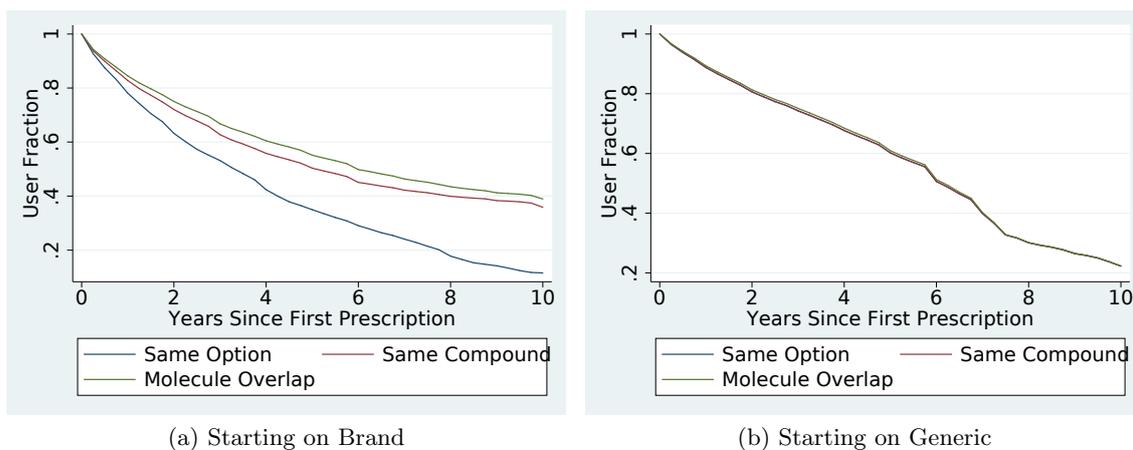


Figure 1: Raw patient-level statistics from the MarketScan dataset, showing a patient's choices over time in relation to his or her initial choice of drug. The x-axis represents years since a user starts on cholesterol therapy, and the y-axis represents the fraction of users that make a given type of choice. In Panel (a), I show whether a user chooses the same branded drug as their initial choice (blue), same molecule (includes generic versions; red), and same family of drugs (includes combination drugs; green). In Panel (b), I show the same outcomes but for users starting on generic drugs. Both figures show significant persistence in choice over time.

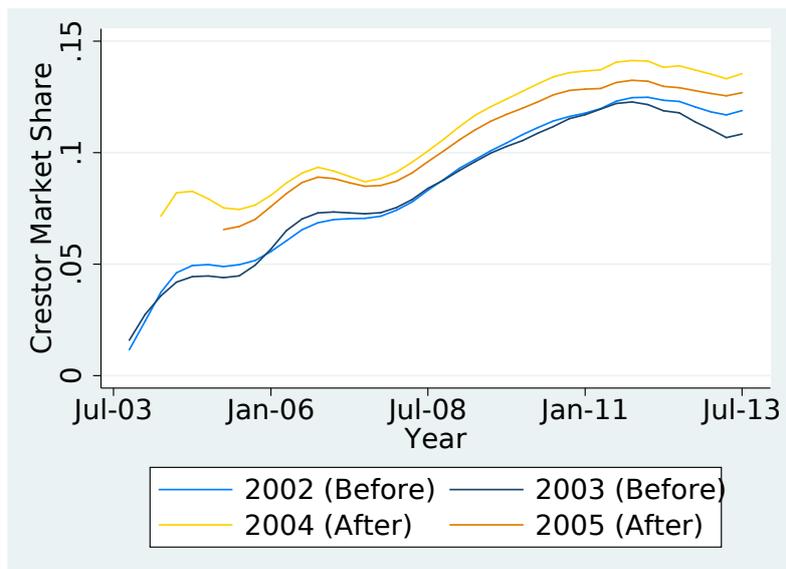


Figure 2: Smoothed Crestor user market share in each quarter by starting-year cohort. The cohort is defined as the year in which a patient begins to take cholesterol medication. The 2003 and 2004 cohorts begin after Crestor is launched (in late August 2003), whereas the previous cohorts start when Crestor is mostly unavailable. The significant difference in this simple case study points to the possibility that starting conditions have an impact on subsequent choice.

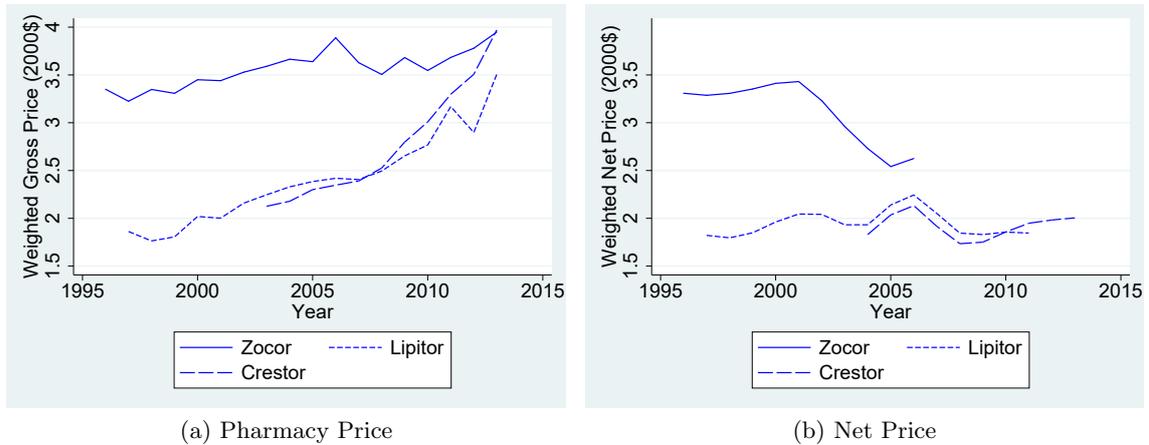
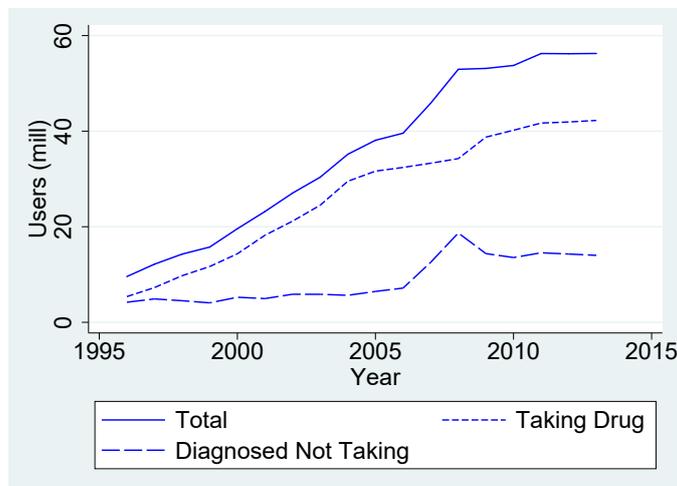
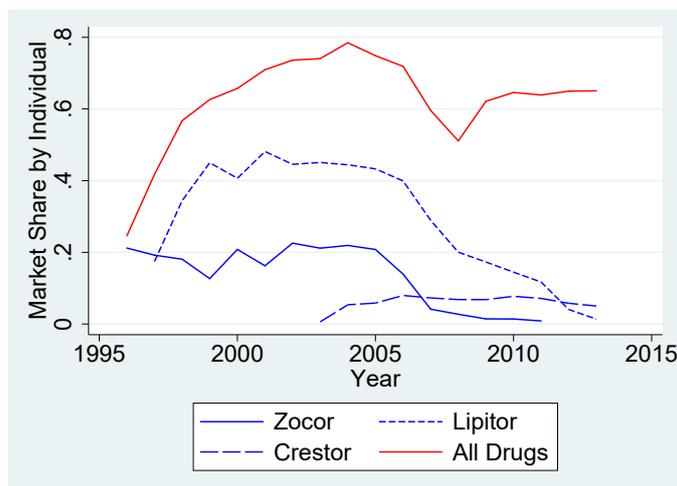


Figure 3: A comparison of average pharmacy and net prices in each year for the three major drugs over the analysis period, showing very different patterns when accounting for rebates. Prices are weighted across medium-intensity treatment level dosages for each drug, with the pharmacy prices coming from median prices in the MarketScan claims data. Smoothed net prices incorporate estimated discounts described in Appendix A.2 and are smoothed out to account for unevenness in wholesale purchasing across years. All prices are adjusted to Year 2000 US dollars.



(a) Market Size



(b) Medium-Intensity Market Shares

Figure 4: Graphs summarizing demand in the medium-intensity treatment level anti-cholesterol market (see Section 2.4 for a detailed discussion). *Top*: number of patients who are taking cholesterol medication and who are diagnosed with high cholesterol but not taking medication, based on MEPS data. The sum represents the market size, which grows at a slowing rate over time. *Middle*: given the market size, a plot of the fraction of the market on each branded medication. The red sums together people taking the three main drugs plus people taking generics and other minor drugs. Generic entry significantly reduced market shares of branded Zocor (June 2006 generic entry) and branded Lipitor (December 2011 generic entry).

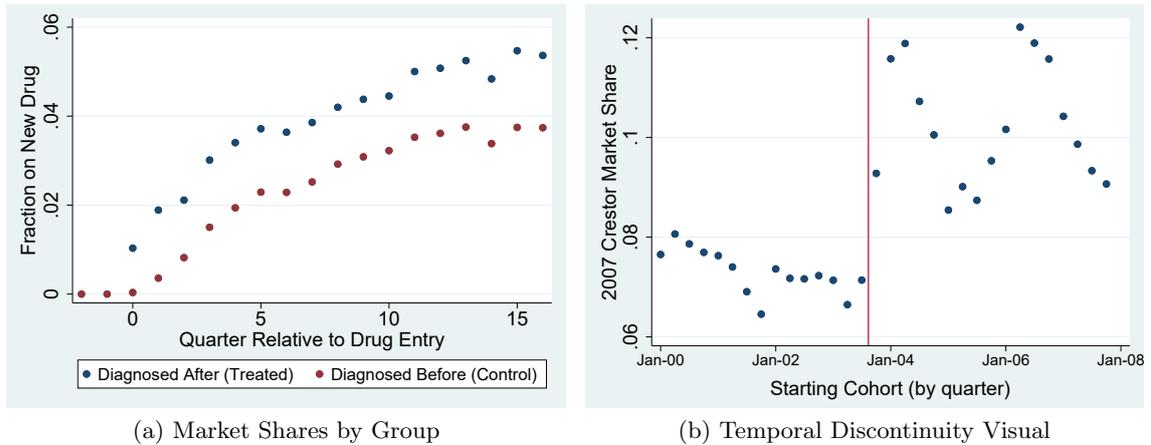


Figure 5: The graph on the left compares the market shares for the new drug across “before” and “after” groups, pooling all branded drug entry events in the sample period. The mean in each period is over the people still in the MarketScan dataset and still on cholesterol medication, generating non-monotonicities. The graph on the right highlights the temporal discontinuity aspect of the research design. It plots 2007 Crestor market share by cohort, where cohorts are defined by the quarter in which a patient starts on anti-cholesterol medication. The vertical line represents the quarter in which Crestor is launched, so the cohorts to the right have an opportunity to select Crestor as the initial treatment.

	t+3	t+6	t+9	t+12	t+15
Initially Chose New Drug	0.687*** (0.00922)	0.632*** (0.0142)	0.547*** (0.0171)	0.580*** (0.0196)	0.545*** (0.0197)
Observations	1354497	1005048	774485	611098	419943
R^2	0.424	0.230	0.159	0.126	0.095

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2: Estimates of Equation (1) for 3, 6, 9, 12, and 15 quarters after drug launch, using the full unbalanced panel of MarketScan enrollees. The coefficients reflect the effect of randomly assigning a user to a drug on the probability that the user chooses the same drug in later quarters.

	t+3	t+6	t+9	t+12	t+15
Initially Chose New Drug	0.685*** (0.0218)	0.542*** (0.0317)	0.473*** (0.0344)	0.408*** (0.0368)	0.380*** (0.0363)
Observations	140133	127515	121390	117554	101802
R^2	0.420	0.205	0.141	0.107	0.088

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3: Estimates of Equation (1) for 3, 6, 9, 12, and 15 quarters after drug launch, using the balanced panel of MarketScan enrollees from 2003 to 2012. The coefficients reflect the effect of randomly assigning a user to a drug on the probability that the user chooses the same drug in later quarters.

	t+3	t+6	t+9	t+12	t+15
Initially Chose Drug	0.543*** (0.0593)	0.571*** (0.0836)	0.600*** (0.101)	0.375** (0.123)	0.463*** (0.121)
Observations	81599	59061	45697	35076	23573
R^2	0.257	0.138	0.101	0.065	0.064

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4: Estimates of Equation (1) for 3, 6, 9, 12, and 15 quarters after drug launch, in a sub-sample of patients who start treatment after a hospitalization. The coefficients are comparable but noisier relative to those in Table 2.

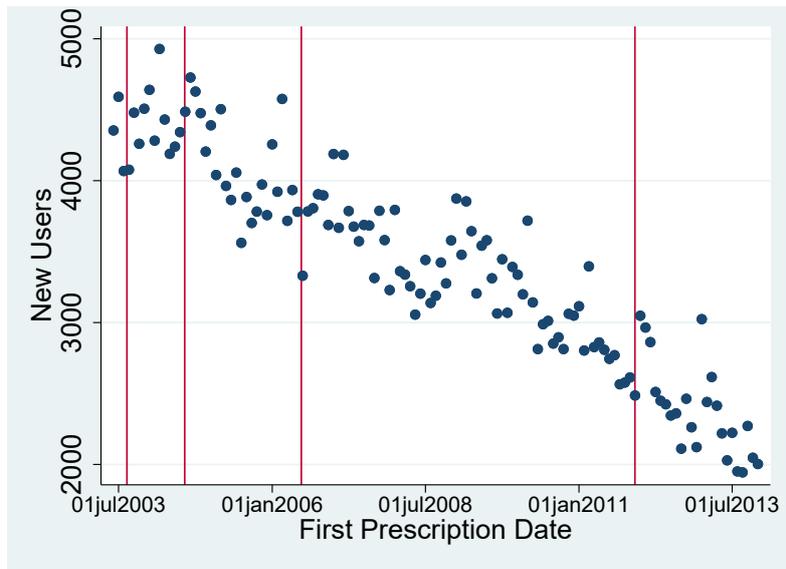
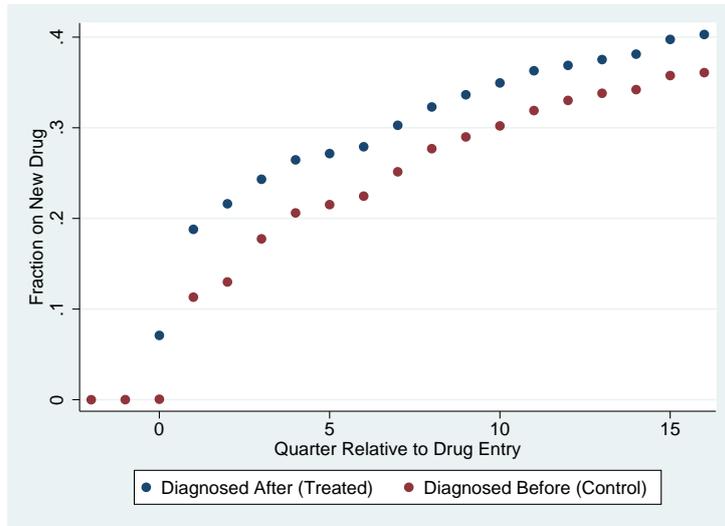
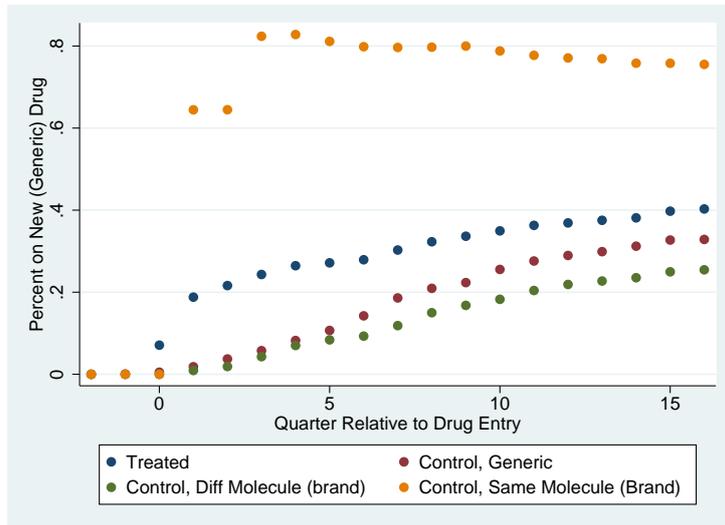


Figure 6: A plot of the number of new patients starting on cholesterol medication during each month in the balanced 10-year MarketScan panel. The vertical lines represent the entry dates of Crestor, Vytorin, generic Zocor, and generic Lipitor. The graph generally shows a smooth trend, with no noticeable spikes around entry events.



(a) Generic Zocor



(b) Generic Zocor Breakdown

Figure 7: Shows a breakdown of the fraction of patients who choose generic Zocor by (a) treatment and control groups defined using the 180-day window used for Figure 5 and (b) treatment and several control groups, with the control groups broken down by type of initial drug choices. The first graph shows a similar pattern to those for branded drug entries, with those in the treatment group more likely to adopt generics. The second graph breaks the controls into patients who started on branded Zocor (orange), patients who started on a different drug (green; typically Lipitor and Crestor), and the few patients who start on generic Mevacor (red). The graph shows that over 80% of users who start on branded Zocor switch to the generic version, whereas the adoption rate of Zocor generics is much lower for the other groups. This is suggestive of greater switching costs across molecules.

Figure 8: A plot of market share (prescriptions) by medical school graduation cohort in the Medicare Part D data. The graph shows little variation across graduation cohorts, suggesting that doctors are only minimally influenced by their own past choices. The only trend is an increase in adoption of generic drugs in recent medical school cohorts

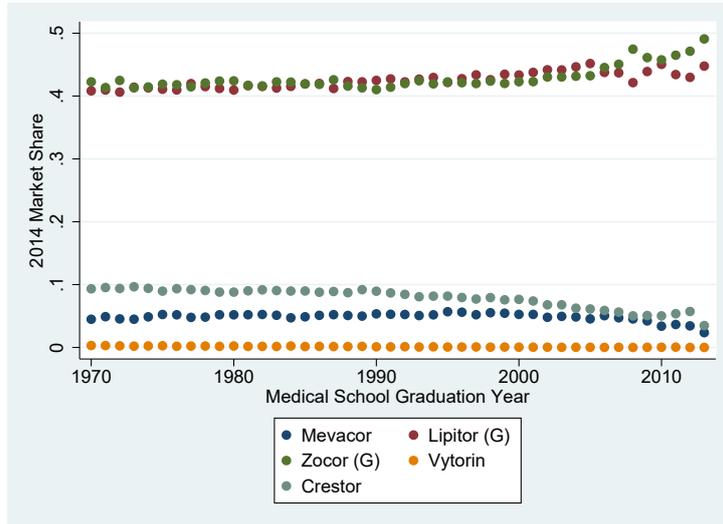
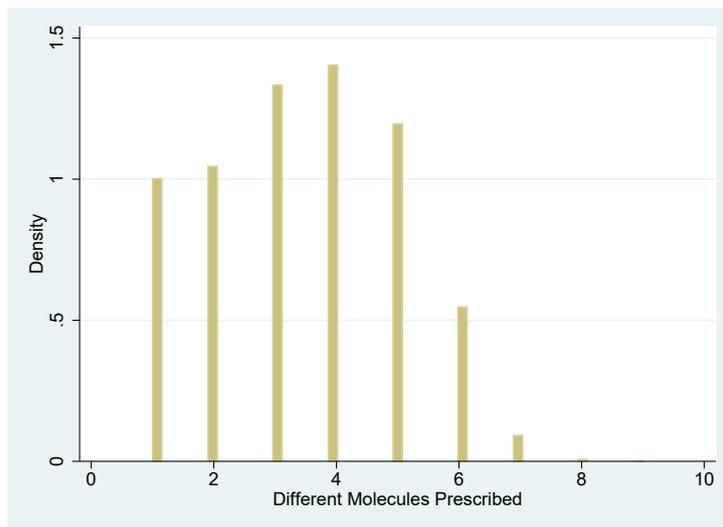


Figure 9: A histogram showing the variety in number of distinct molecules prescribed (counting generic and brand of the same molecule as one drug), across the 210k doctors that prescribe any cholesterol medication to Medicare Part D patients in 2014. About 75% of doctors prescribe at least three drugs with different molecules.



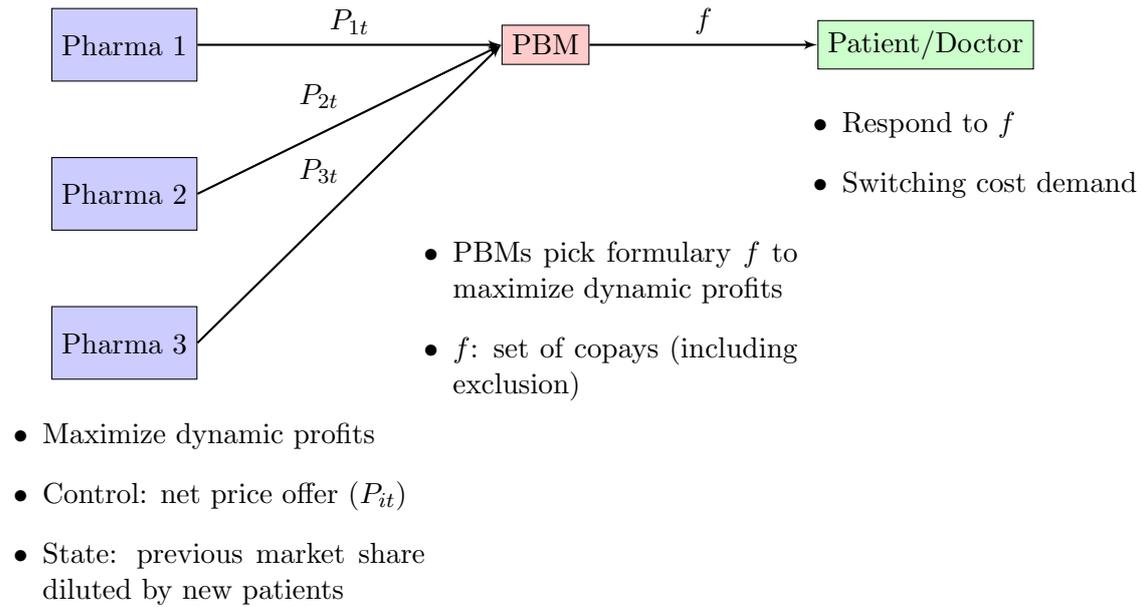


Figure 10: A diagram summarizing the pricing structure in chronic drug markets with inertia, as discussed in detail Section 4. The key interactions in the model are: i) drug companies make take-it-or-leave-it net price offers to a representative PBM that maximize dynamic profits ii) the PBM picks a formulary to maximize its (possibly dynamic) profits iii) patients make drug choices based on the formulary picked by the PBM, while also factoring in switching costs and drug quality. Drug companies and PBMs take into account patient behavior, in particular history-dependence, when making decisions. The dynamic game has state variables equal to the previous period's market share, diluted by the new patients, who are unattached to any drug. The drug companies use net price bids as their control.

Table 5: Estimates of Equation (2). The results show an inelastic response to cost sharing and significant inertia. Incorporating the instrument and the term for previous intensive margin choice diminishes some of the raw estimates. The instrumental variables specifications (2SPS and 2SRI) have standard errors clustered at the cohort-level to reflect the source of the exogenous variation used in the estimates.

	(1)	(2)	2SPS-1	2SRI-1	2SPS-2	2SRI-2
Copay (α)	0.662*** (0.00225)	0.719*** (0.00253)	0.633*** (0.101)	0.724*** (0.141)	0.668*** (0.111)	0.739*** (0.149)
Same Drug (γ)		3.341*** (0.00235)	2.674*** (0.567)	2.502** (0.778)	1.460*** (0.204)	2.070*** (0.277)
Prev. Drug (ν)					1.385** (0.494)	0.330 (0.494)
Drug FE	x	x	x	x	x	x
N	49,413,573	49,413,573	49,413,446	49,413,446	49,413,446	49,413,446

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

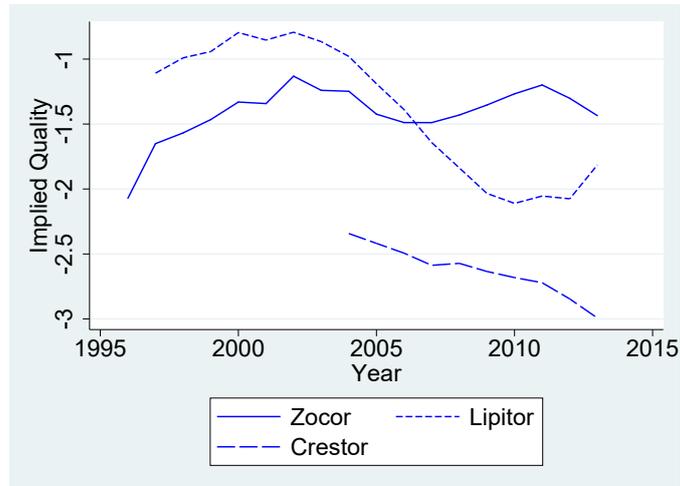


Figure 11: A plot of the δ_{jkt} coefficients from Equation (2), with generic Zocor and generic Lipitor qualities replacing their branded counterparts in the series once they enter. The data shows that Zocor has lower implied quality than Lipitor until its generic version enters. Both Lipitor and Crestor experience declines in implied quality once generic Zocor enters. This is because the fraction of diagnosed patients taking a drug stays steady, but generic Zocor begins to take market share from its competitors. Other quality fluctuations are driven by advertising campaigns and changes in medical evidence.

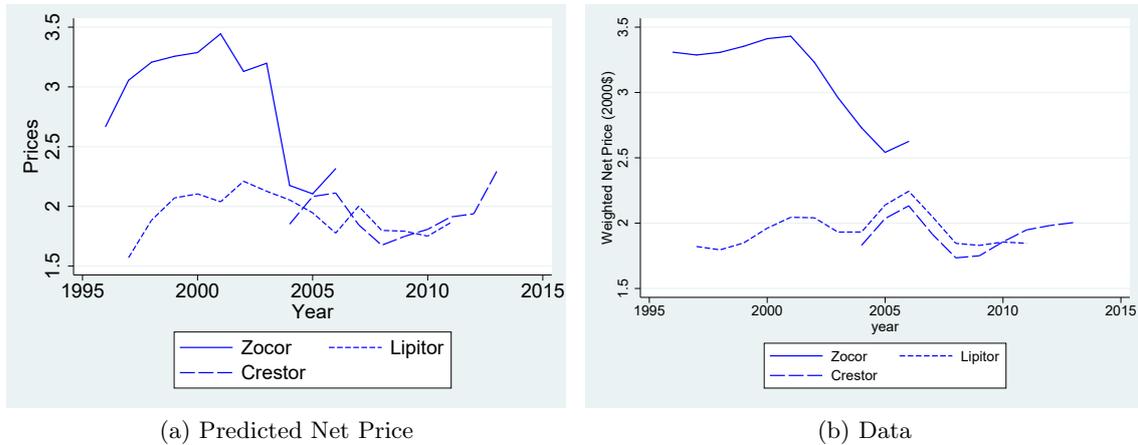
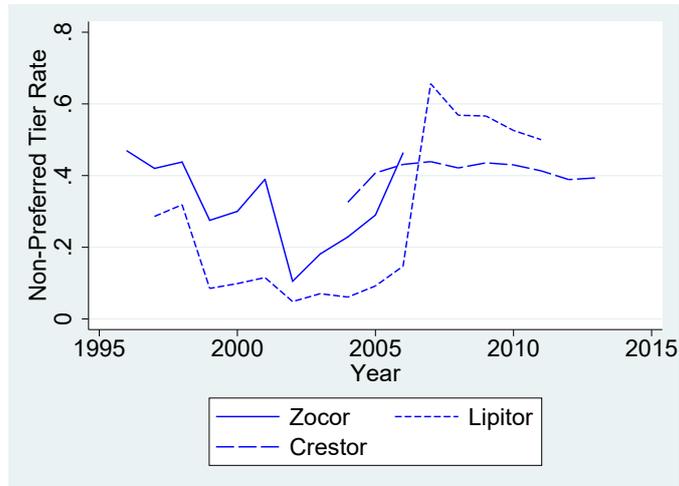


Figure 12: A comparison of model predictions and data on net prices. The model captures the general dynamics in the net price data, including the key features laid out in Section 2.7. The model rationalizes the gap between Zocor and Lipitor prices in early years, the decrease in net prices after generic Zocor enters the market, and the general increase in price in anticipation of generic Zocor and Lipitor entry. Aspects not matched are due to inflexible features of the model discussed in greater detail in Section 5.3. This includes the size of Zocor’s drop in net prices, Lipitor’s pricing in 2006 relative to 2005, and pricing in 1996 and 1997.

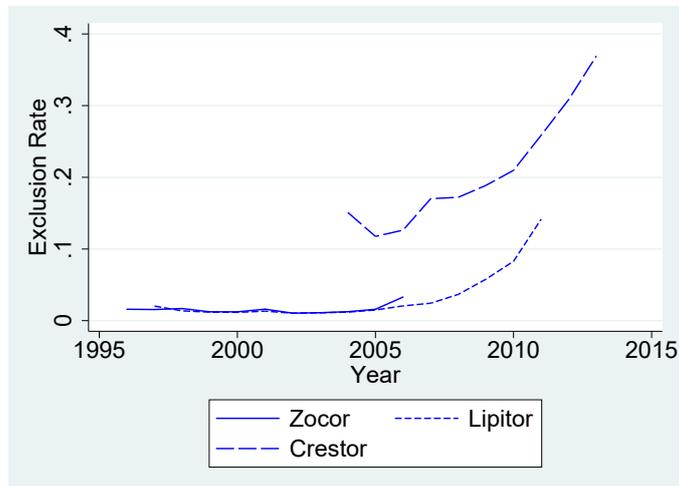
Parameter	Estimate
θ	0.973 (0.083)
χ	0.826 (0.110)
β^{PBM}	0.100 (0.109)
Merck Boost (b_1)	1.155 (0.127)
Generic Boost (b_g)	0.449 (0.360)
κ	0.086 (0.130)

Standard errors in parentheses

Table 6: Parameter estimates from the model described in Section 4. Standard errors are computed using the standard maximum likelihood framework.



(a) Non-Preferred Tier



(b) Exclusion Probability

Figure 13: Graphs capturing the equilibrium formulary distribution predicted by the model. The first plots the share of formularies in which each drug is in the non-preferred tier and the second plots the rate of exclusion for each drug. The predictions generally reflect trends found in the MarketScan data and industry reports: i) no exclusion in the early periods ii) Zocor is in the non-preferred tier at higher rates relative to Lipitor iii) after generic Zocor enters, higher copays and formulary restrictions are put on the remaining drugs.

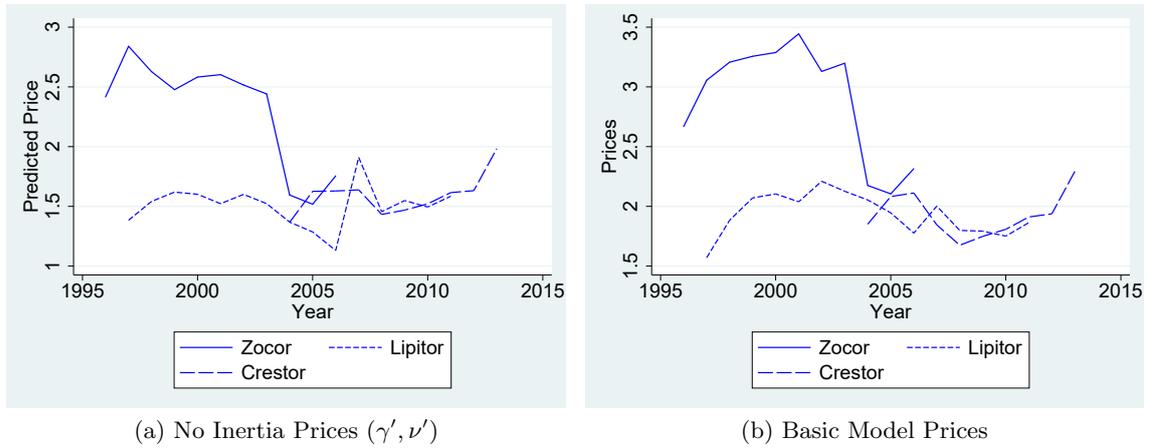


Figure 14: Counterfactual prices under a scenario with no drug-specific inertia ($\gamma' = 0, \nu' = \nu + \gamma$). The price levels are generally lower in all periods, and prices exhibit general deflation over time as competition becomes stronger in each period.

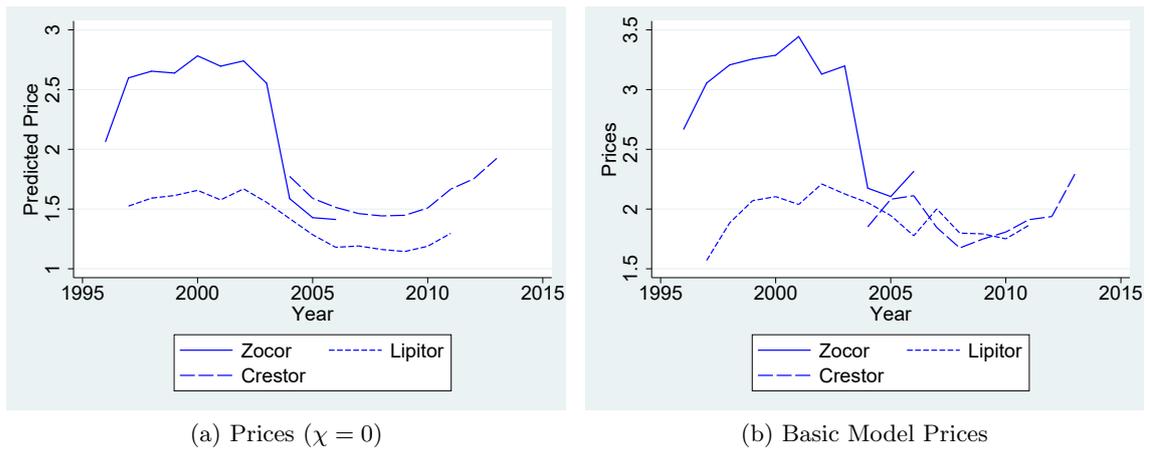


Figure 15: Counterfactual prices under a scenario with no additional penalty for excluding popular drugs ($\chi = 0$). The relative prices remain similar, but price levels are much lower and prices trend downwards over time as more competition enters.

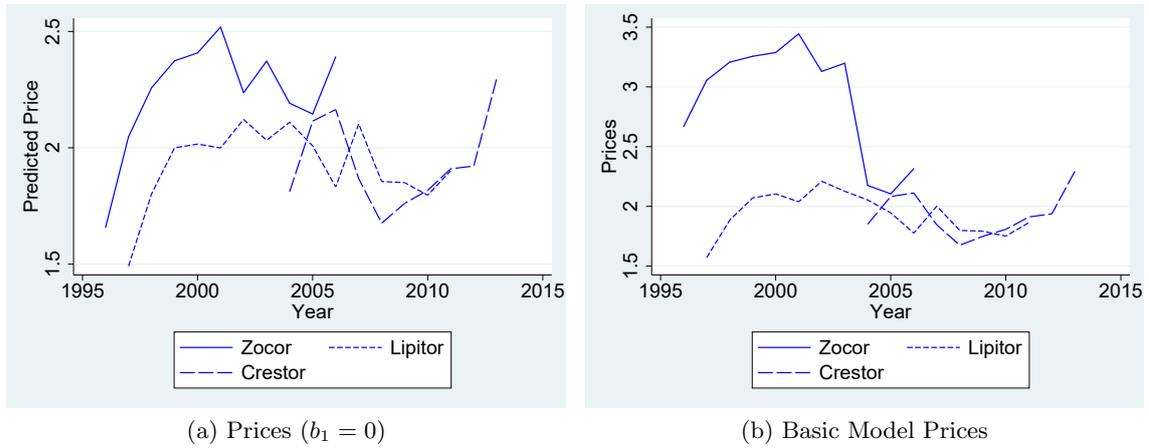


Figure 16: Counterfactual prices under a scenario with no Merck-Medco relationship ($b_1 = 0$). The qualitative result that Zocor is still priced higher relative to Lipitor is maintained in the counterfactual, highlighting the impact of the discrete copay structure on . Other qualitative aspects of the price dynamics also remain similar.

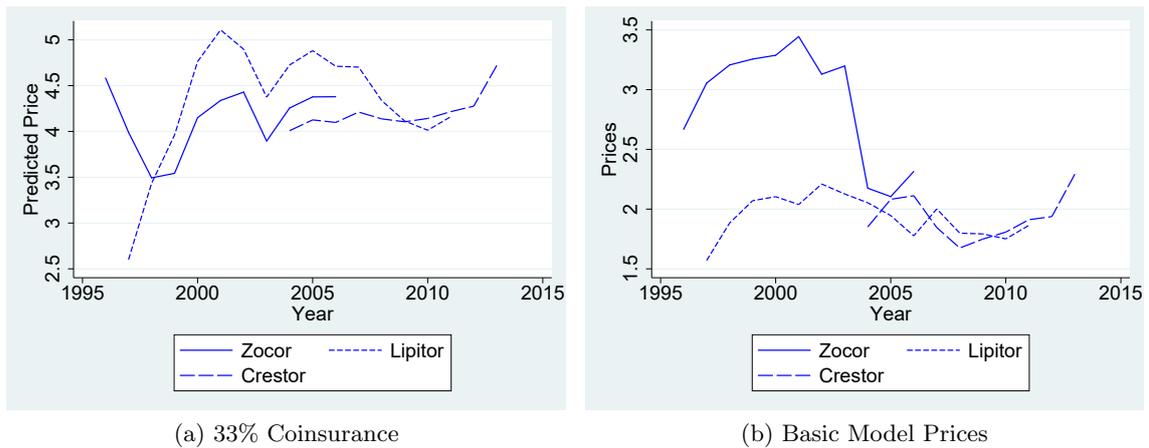


Figure 17: Counterfactual prices under a scenario where patients face a fixed fraction of the prices set by drug companies, outlined in Section 6.3. The results show much larger swings in prices, particularly for Pfizer as it enters and gains market share. Pfizer also ends up pricing higher than Zocor for much of the period, a qualitative departure from the actual data.

Scenario	Average (\$)			Net Revenue (\$ bill)
	Zocor	Lipitor	Crestor	
Baseline	2.83	2.03	1.91	151
<i>Varying Model Parameters:</i>				
No Merck-Medco	2.27	1.98	1.91	141
No inertia	2.17	1.47	1.57	119
No exclusion penalty	2.13	1.41	1.58	120
<i>Pricing Structure Changes:</i>				
30% Coinsurance	4.37	4.53	4.18	195
Single Tier	2.96	1.96	1.77	158
Commit one preferred	2.71	1.93	1.77	139
Commit one exclusive	2.03	1.00	1.40	80
No insurance	1.40	1.49	1.39	80

Table 7: Statistics summarizing the results from the counterfactual exercises. i) Baseline: Best fit model. ii) No Merck-Medco: no relationship between the two, sets $b_1 = 0$ iii) No inertia: no drug-specific inertia, with effects shifted to preferences for any drug ($\gamma' = 0, \nu' = \nu + \gamma$). iv) No exclusion penalty: PBM does not have extra aversion to exclusions $\chi = 0$. v) 30% Coinsurance: patients face coinsurance rate of 30% with no PBM involvement. vi) Single Tier: PBMs can only choose inclusion or exclusion for branded drugs, with included drugs having copay equal to the preferred copay level. vii) Commit one preferred: PBM commits to only picking formularies that have one branded drug in the preferred copay tier. viii) Commit one exclusive: PBM commits to formularies that include at most one branded drug. ix) No insurance: patients face full price, assuming same α coefficient. Average prices are computed in a weighted manner across the 18 years from 1996-2013, are scaled to daily prices, and only include branded prescriptions. Total payments is computed by multiplying computed equilibrium net prices by market share and by market size, and summing over all years. Net revenue captures revenue to drug companies over the 18 year period, and excludes PBMs profits.

Appendix

A Additional Data Background

A.1 MarketScan User Panels - Unbalanced and Balanced

Given the significant entry and attrition of users in the MarketScan data, I work with both an unbalanced and a balanced panel of users in my analysis. Each has its advantages, which I factor in to my later analyses.

The unbalanced panel encompasses all users in the Truven dataset, which yields great statistical power but raises issues surrounding user history and weighting. Using the raw data is the simplest approach and provides the greatest statistical power for my analysis, especially in the 2000s, when the dataset covers a sizable fraction of the US population. However, the rapid growth in people covered creates issues in terms of weighting and composition. For example, if we compute average copays or prices for all users in the dataset, changes could just be driven by the nature of entering users rather than changes in price facing existing users. Similarly, when we pool all entry events to study history-dependence, events later on in the period will have a higher weight.

To alleviate some of these problems, I also construct a smaller balanced panel from 2003-2012. I choose the period based on the 10-year period that gives us the most number of users, which turns out to be around 2.5 million. Starting any earlier is problematic, because most users from the 1990s have exited the data by the end of the coverage period. Of course, a balanced panel selects on specific types of employers that consistently use Truven services and on employees who tend to stay at the same firm over time.

A.2 Constructing Net Price Estimates

In this section, I provide a more detailed description of the issues in constructing net price estimates, the data I collect, and my methodology for dealing with problematic aspects of the data.

A major challenge here is that net sales estimates are at the drug level, but different dosages of the same drug are sometimes priced differently. For example, a 40MG Lipitor tablet has a gross price that is \$1 more expensive than the more commonly used 10MG tablet.¹⁰³ Since I am focusing on the aggregate patterns in the medium intensity dosages, I want to avoid interpreting shifts in other dosage levels or in form of purchase as changes in rebate levels, so I construct overall weighted net and gross prices, and assume that rebates are roughly constant at the drug level.

To construct my estimates, I collect revenue data from SEC filings and IMS Top-Line data, US aggregate demand estimates from the Medical Expenditure Panel Survey (MEPS), and additional net price estimates from SSR Health. I then construct net prices based on the data available in each year.

The financial filing data comes from annual company SEC filings (10-K) and conference call transcripts, between 1996-2013. At the start of the period, companies generally

¹⁰³Crestor actually sets the same price for all dosage levels.

report global sales by disease area, but progressively offer more granular data, with the current standard being a breakdown by drug and geographic area (US, Europe, global). Roughly, the financial data coverage is as follows:

1. 1996-2000: Merck reports global aggregate sales for its cholesterol drugs (Mevacor and Zocor) in its 10-K filings; Pfizer reports global Lipitor sales in its 10-K filings
2. 2001-2003: Merck begins to report Zocor global sales separately in 10-K filings
3. 2004-2006: Merck reports US quarterly net sales for Zocor in its conference calls
4. 2003-present: Merck (at the time Schering-Plough) reports global sales of Zetia and Vytorin.
5. 2003-present: Pfizer and Crestor begin reporting net US sales in its 10-K filings.

In addition, I collect additional net revenue data from IMS Top-Line Market Data¹⁰⁴ and from media reports on the cholesterol market. The IMS Top-Line data is currently available from 2007

In terms of demand data, I construct estimates of average demand from MEPS. I break this down by different dosage levels and seller (pharmacy vs. mail order) for each of the cholesterol drugs. MEPS covers the entire period from 1996-2013.

Next, I focus on gross prices. To do this, I 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 the same across all entries with the same NDC number, while the “pay” variable partially reflects pharmacy markups.

Finally, I obtained estimates of net sales and quantity sold from SSR Health, a pharmaceutical consulting company. They estimate net prices for a large number of drugs since 2007 at the quarterly level, and shared with me their estimates for the cholesterol market. I use their numbers for the post-2007 period, but my estimates using raw financial filing data also lines up with their estimates.

After collecting this data and aggregating it to an annual level, I take the following procedure to calculate net prices:

- In years with reported net US sales, I just divide net sales by estimated US quantity from MEPS.
- In years where only net global sales are reported, I arrive at a net US sales number by making the assumption that all of the global rebates are from rebates paid in the US. I verify this assumption in years with both net global and net US sales numbers.

$$\text{Net US Sales} = \text{Net Global Sales} - (\text{Gross Global Sales, IMS} - \text{Gross US Sales, IMS})$$

¹⁰⁴<http://www.imshealth.com/en/about-us/news/top-line-market-data>

I also make the following corrections to account for issues created by aggregating to an annual level:

- Zocor generics enter in June 2006. Using aggregate quantity for the whole year creates a misleadingly low price for pre-generic net price. This is because after generics enter, branded drugs typically price at a duopoly level to compete with the exclusive generic.¹⁰⁵ Instead, I estimate the fraction of units sold in the first two quarters, and use quarterly revenue data to arrive at a pre-generic net price.

To arrive at a discount per pill, I compare the overall net price estimates to overall gross sales estimates, both by looking at IMS gross sales numbers when available and by constructing my own estimate of gross sales using pharmacy values. The pharmacy payment should roughly be an upperbound on the true net sales number, as illustrated in Figure.

$$\text{Avg. Gross Price} = \frac{\sum_{ij} P_{ij} X_{ij}}{\sum_{ij} X_{ij}}$$

where i is the drug and j is the seller (pharmacy or mail order), X is the sales estimate from MEPS and P is either the AWP or “pay” variable from MarketScan.

B Additional Reduced Form Results

C Robustness - Methodology, Other Events, and Other Chronic Drug Markets

Next, I turn to the issue of the robustness of the results. My estimates are not particularly sensitive to the window I use and the simplification to quarterly data. In addition, the same inertial patterns are present around negative clinical trial announcements and in other chronic drug markets, suggesting that inertia is an important force in many drug-related settings.

One concern is the robustness of the results to the window around the launch events, which I address by using a narrower window to select treatment and control users. As mentioned earlier, the main motivation for selecting users starting at a similar time is to avoid picking up on any changes over time in diagnosis guidelines that may shift the composition of new users. In Table 8, I report the same coefficients when selecting users based on 60-day windows around the launch date. The coefficients are again noisier, as our sample is cut by a third, but generally similar to the ones using 180-day windows. The increase in coefficients after period 9 is influenced by the composition of drug entry events and their associated user attrition rates.¹⁰⁶

I also find little evidence that existing patients try new drugs once and switch back to their old drug, a behavior that I might be attributing to inertia if quarterly aggregates

¹⁰⁵The first generic entrant typically has an 180-day exclusivity period based on provisions in the Hatch-Waxman Act.

¹⁰⁶For example, there was very high attrition in 2005, so Crestor and Vytorin would dominate the early coefficients, but many users for that event drop out, leaving other events to influence the coefficient.

obscure fast switches. Under a learning model, it's possible that the control group may try the new drug at the same rate as the treatment group, but then switch back to what they started on. I verify that this is not the case by re-running the analysis with the outcome "prescribed new drug in current or previous period," which can also be thought of as an "ever tried" variable. To make sure that my quarterly simplification does not miss users who try the new drug once, I generate the variable using raw claims data. This ends up not being a large issue, as it only affects the data in 0.3% of all user-quarters, and, unsurprisingly, the estimates are almost identical to those in Table 2.

Next, I turn to other types of variation in starting conditions beyond drug availability, focusing on starting cohorts around the announcement of influential clinical trials. In the anti-cholesterol market, drug companies routinely run post-approval clinical trials to ascertain long-term health outcomes, particularly in chronic disease areas. Some clinical trials are influential in changing beliefs about drug quality, and if inertia is present, one might expect some differences in cohorts who start before or after a big announcement. I test this by looking at the announcement of the results of the ENHANCE trial in January 2008,¹⁰⁷ which found that the combination drug Vytorin did not reduce heart attacks more than statins like Zocor, which had gone generic in 2006. Prior to the study, Vytorin was thought to be more effective, but the results suggested that they were a much less cost-effective option. Figure 22 shows the market share of Vytorin in cohorts starting treatment around the announcement of the trial. Consistent with the core results, the "before" cohort has a steadily higher Vytorin market share relative to the "after."

Finally, Figure 21 in the Appendix shows that these patterns hold to different degrees in other major chronic disease areas. The graphs report results from applying my methodology to other areas, and suggest that history-dependent demand may also play a significant role in the multiple sclerosis, diabetes, and chronic obstructive pulmonary disease (COPD) markets. The main caveat here is that these markets are more medically complex, which brings into question my identification assumption. Some drugs may target patients for whom the existing therapies are inappropriate, so many patients may wait for the new drug to launch before starting, violating my identification assumption.

C.1 Heterogeneity by User Characteristics

Next, I turn to documenting heterogeneity in history-dependence across user characteristics such as age, gender, and observable health characteristics, which will be important for demand modeling and gaining additional insight into microfoundations, given previous findings about demographic predictors of risk aversion and adoption.

My approach is exactly the same as the instrumental variables one above, except now I add in interactions. What I'm doing in effect is to separate out the same analysis for different groups of patients. For example, we would be looking at difference in female market share after 4 years, scaled by the initial difference.

The first set of characteristics I look at are purely descriptive, based on user demo-

¹⁰⁷<https://www.wsj.com/articles/SB120052963183895929>

graphics including age and gender. I find small negative effects for both age and gender in terms of history-dependence at all points in time. In other words, women and older individuals are more likely to switch than their counterparts. This may be surprising given other results about the correlation of age and gender with risk aversion and adoption in the literature. However, it is possible that these characteristics are correlated with unobserved health state or experience with generics.

The second set of characteristics are based on health at the time of first prescription, which can lend insight into the micro-foundations. I find coefficients suggesting that healthier individuals display much lower levels of history-dependence. For each additional drug a patient is on at the time they start anti-cholesterol medication, the causal effect up by 10 percentage points. Similarly, patients starting treatment during an inpatient episode exhibit about a 30 percentage point increase in causal effect.

Finally, I put all of these factors together in a horse race regression in Table ??, finding that the health characteristics appear to be the main drivers, and that the female and age coefficients are now very weakly positive, more in line with expectations given a risk aversion-based story of history-dependence.

C.2 Using the Methodology to Compare Long-Term Health and Spending Impacts

An interesting byproduct of my empirical strategy is a way to estimate the long-term treatment effects of a given drug, relative to existing treatments. I discuss the approach, including the assumptions and conditions needed for it to work, and offer a proof of concept.

The approach is to use the same treatment and control groups as in my main methodology, but to compare average health outcomes instead of usage across the two groups. The intuition behind the approach is that if users are randomly assigned to a drug, and exhibit a high degree of inertia, any health outcome differences should be attributable to long-term usage of different drugs. One can use the same instrumental variables framework as in Equation (1), but alter it to be

$$H_{it} = \alpha + \gamma_t S_{it} + \epsilon_{it} \quad (4)$$

where H represents health outcomes including hospitalization, health care spending, and common side-effects. S can be some indicator of usage of the new drug between period 0 and period t , such as continual usage or quarters of use.¹⁰⁸ Once again, we would instrument for S_{it} using the before or after treatment variable Z_i .

The interpretation of the coefficient would be the effect of taking the new medication over the existing set of medications. For example, if one were analyzing the entry of Crestor, γ would measure the health differences generated by Crestor relative to the existing standards of treatment. This is again a treatment-on-the-treated estimate, normalizing the average difference by the rate of adoption. Another issue here is that there could be local average treatment effects at play, as the complier group may be

¹⁰⁸One might use “uninterrupted usage” as the predictor if there may be non-linear effects, whereas one would use “quarters of use” if benefits are linear in usage.

especially well-matched to the entering medication. In this case, the coefficient should be interpreted as the value of the new treatment for the complier population.

As a proof of concept, I analyze the impact of new drugs on hospitalizations using this framework, and find very little difference in outcomes. Hospitalizations are a key health outcome in many settings, as reducing occurrences can lead to large savings.¹⁰⁹ In Figure 27, I plot hospitalization rates for the treatment and control groups associated with Crestor entry, and find very little long-term differences between the two groups.¹¹⁰ Similar analysis of other drug entries in this period yield similar results.¹¹¹ This results is not too surprising, as statins are generally very similar in efficacy and side-effect profiles.

The framework can also be used to assess spending, the other half of the cost-effectiveness equation, and I find more interesting dynamic spending differences that point to possible long-run cost-reducing strategies on the part of PBMs or payers. For my basic analysis, I just plot mean annual spending on cholesterol medication by group.¹¹² I find that newer branded entries generally lead to persistently lower annual costs for the treated group, which is unsurprising, given that later drugs tend to have lower prices. One also finds interesting patterns in spending for cohorts around generic entry events. Figure 28 in the Appendix presents two entry events with significant and time-varying differences across cohorts in spending. For the entry of Livalo, a less popular statin launched in 2010, we see an initially lower spending level for the treatment group, which then jumps above the control group when generic Lipitor enters and most Lipitor users move to the cheaper generic. Therefore, it may be long-run cost-effective to incentivize users to stay on Lipitor. A similar story emerges by analyzing groups of users around the entry of generic Zocor, which I show in the second graph in Figure 28. Here, I find that the patients in the control group who picked a different branded drug (mostly Lipitor and Crestor) end up spending more over time. Control patients who pick branded Zocor start off spending more, but converge to a much lower price in the long-run.

C.3 Micro-Foundations - A Detailed Discussion

Here, I discuss the possible micro-foundations driving the history-dependence uncovered above. There are two general classes of explanations for the results, rational and behavioral, both of which have support in the data. In addition, there is the issue of the role played by doctors, which is important in thinking about how much market position carries over when the market size is growing and how large inertia may be in markets for acute medication.

There is some support in the data that the observed inertia is generated by experience good effects. Experience good models typically involve a risk-averse user learning about

¹⁰⁹I also look at differences in general health spending as an outcome.

¹¹⁰The initial difference comes from the fact that many new users start after a hospitalization, so the control group has a spike about two quarters before the treatment group. The results are similar when I restrict to heart-related hospitalizations.

¹¹¹The one surprise is that the results also hold for Baycol, a drug that was later pulled from the market due to safety concerns.

¹¹²For now, I use pharmacy prices in the MarketScan dataset.

the qualities of products through consuming them. The higher the risk aversion, the lower the likelihood of switching. In the Appendix, I show that heterogeneity in inertia mainly stems from differences in health state, with unhealthy patients and those on many other drugs much less likely to switch. These patients are probably more risk averse in their decisionmaking, bolstering the case for experience good effects.

As in many dynamic individual behavior settings, habit formation can also explain inertia. Once a user develops a routine of which drug to take, they may find that repeated usage brings them higher utility. This explanation is harder to test, although the additional evidence I provide on the relationship of switching to duration of usage does suggest this may be driving part of the effects.

Another explanation of the results is the presence of physical switching costs driven by health considerations. Physical costs are less obvious in the anti-cholesterol market setting, where treatments are fairly uniform. However, there is some discussion and research in the medical literature on poor compliance in patients after switching, which is a form of switching cost. In addition, patients need to go to their doctors to obtain a change in prescription, another form of switching cost. In other areas, such as multiple sclerosis, there are clearer medical costs to switching. Patients need to go through a tapering of the old treatment and a waiting period before starting on the new one.

Behavioral explanations for inertia are led by inattention. One possibility is that patients do not pay attention to prices and new options, and stick with their existing option unless put into an active choice environment. In refinements to my basic demand estimation, I find that patients are more likely to switch in periods where they are hospitalized, providing some evidence that inattention may be at play. However, this is also consistent with the rule of thumb discussed previously. More definitive evidence on this, through exogenous shocks to attention such as exposure to advertising, would open up the possibility of informational interventions in changing observed inertia.

Other behavioral phenomena such as the endowment effect or loss aversion may also be factors in this setting. The main evidence pointing to this is the product hopping phenomena discussed earlier. People appear to switch from incremental products to generics of the base product and vice versa, suggesting that they dislike either having to pay more or to lose out on quality.

D Illustrative Model - Inertia and Copay Tiers

In this section, I provide a simple model to capture the qualitative effects from the interaction between patient inertia, PBMs, and the copay tier structure. The key implication is that entrants will either price close to existing competitors, or choose to undercut the incumbent by a significant amount to push them into the higher copay tier.

D.1 Model Outline

The simplified model contains the following features:

1. There are two firms. Firm 1 has an established drug with price p_1 (exogenous for

now) that is being taken by γ of the current patient pool (normalized to have unit mass).

2. Firm 2 has an entering drug that is identical to Firm 1's drug. To begin with, I assume that Firm 2 is the only active price setter.
3. The intermediary (PBM) takes p_1, p_2 , and assigns each to a copay tier (no possibility of exclusion).
4. The formulary arrangements are based on two copay tiers: \underline{c}, \bar{c} with $\underline{c} < \bar{c}$
5. The intermediary's objective function is summarized as a weighted average between consumer surplus (simplification of demand for PBM services) and expenditures:

$$\max_f W(f) - \theta \sum_{i=1}^2 p_i q_i(f)$$

where $W(f)$ is the value to the consumer (normalize to 1) minus the copay.

6. Demand is assumed to be very simple: patients already on the drug owned by Firm 1 will keep taking it, and new patients will pick the drug with lower copay, unless they are equal, in which case they split evenly across the two drugs.

D.2 PBM Decision

We start with the PBM problem. Assume for now they pick between three formulary arrangements:

$$\mathcal{F} = \{(\underline{c}, \underline{c}), (\underline{c}, \bar{c}), (\bar{c}, \underline{c})\}$$

Given p_1, p_2 , the $(\underline{c}, \underline{c})$ formulary yields:

$$v_{same} = (1 - \underline{c}) - \theta \left[\gamma p_1 + \frac{1}{2}(1 - \gamma)p_1 + \frac{1}{2}(1 - \gamma)p_2 \right]$$

while the (\bar{c}, \underline{c}) arrangement ("2 preferred") yields:

$$v_{2p} = \gamma(1 - \bar{c}) + (1 - \gamma)(1 - \underline{c}) - \theta [\gamma p_1 + (1 - \gamma)p_2]$$

and the (\underline{c}, \bar{c}) arrangement ("1 preferred") yields (everyone picks Firm 1):

$$v_{1p} = (1 - \underline{c}) - \theta p_1$$

D.3 Entering Firm Pricing

Given the PBMs behavior, we then evaluate the incentives facing the entering firm. Firm 2 then sets its price p_2 relative to a fixed p_1 to maximize profits, taking into account the PBMs incentives.

In the simplest case, setting $p_2 > p_1$ will mean that the PBM will pick the “1 preferred arrangement,” as v_{1p} is larger than both of the other options (cheaper vs. “same”; cheaper and higher welfare vs. “2 preferred”). This mean that they will get no profits from this strategy.

At $p_2 = p_1$, the PBM is indifferent between v_{same} and v_{1p} . Firm 2 can price ϵ below to push the PBM to pick the “same” arrangement. One can also modify the model to have a small fraction of users who really like drug 2, thus boosting the welfare under the “same” arrangement.

So the decision then boils down to the price at which Firm 2 can persuade the PBM to choose the preferred tier.

$$\begin{aligned} v_{same} - v_{2p} &= \gamma(1 - \underline{c}) - \gamma(1 - \bar{c}) + \theta \left[\frac{1}{2}(1 - \gamma)p_1 - \frac{1}{2}(1 - \gamma)p_2 \right] \\ &= \gamma(\bar{c} - \underline{c}) + \frac{1}{2}\theta(1 - \gamma)(p_2 - p_1) \end{aligned}$$

Define d^* is the discount $p_1 - p_2$ at which the difference becomes zero. Setting the expression to zero and re-arranging, we get

$$d^* = \frac{2\gamma(\bar{c} - \underline{c})}{\theta(1 - \gamma)}$$

This means that the discount Firm 2 has to offer is:

1. Increasing in $\bar{c} - \underline{c}$: intuitively how unhappy incumbent users would be if they had to pay a higher copay
2. Increasing in γ : intuitively the degree of incumbent advantage
3. Decreasing in θ : how much PBMs are focused on savings vs. keeping consumers happy.

Finally, we can think about Firm 2 profits. At $p_2 = p_1$, they obtain:

$$\pi_{p_2=p_1} = \frac{1}{2}p_1(1 - \gamma)$$

while at $p_2 = p_1 - d^*$, they obtain:

$$\pi_{p_2=p_1-d^*} = (p_1 - d^*)(1 - \gamma)$$

Pricing a little bit below either of those two points is unprofitable, as they obtain the same demand at lower profits per unit.

Therefore, Firm 2 picks the $p_2 = p_1$ strategy if:

$$p_1 < 2d^* = \frac{4\gamma(\bar{c} - \underline{c})}{(1 - \gamma)\theta}$$

The results suggest that there's more room to undercut existing prices if there is a lower copay differential or lower levels of patient inertia.

D.4 Summary

The main takeaway from the toy model is that it is difficult to gain market share through undercutting a competitor with an existing consumer base. PBMs are less likely to want to move a drug with a consumer base to a higher copay tier, or exclude it entirely, as is possible in my full model. Therefore, a company that wants to gain dominant market share will have to undercut competition by a significant amount.

The model also speaks to the entry strategy of firms, who will either price the same or undercut competitors by a large amount. The key mechanism driving this result is the interaction of PBMs with discrete copay tiers and patient inertia. Specifically, the two key features of copay setting are that the formulary tiers are discrete ($\bar{c} - \underline{c} \neq 0$) and that the PBM has the option to put both drugs in the same formulary tier. In conjunction with stickiness, this pushes the Firm 2's pricing to either be equal or to be dramatically different, because Firm 1 has an advantage by having a pool of captured consumers. The PBM is likely to put Firm 1 in the low copay tier, because it has trouble moving people to capture the discounts, and also has to take into account the copay differential they are imposing on the captured consumers. Therefore Firm 2 can choose either to overcome this by setting a much lower price, or choose to split the share of new patients by setting an equal price. If the equal tier setup was not in the set of possible formularies¹¹³, then Firm 2 would have more of an incentive to price lower, as pricing the same as Firm 1 will push them into the higher copay tier.

The prescription drug market is different to other markets with intermediaries, because in other contexts such as supermarket shelf space and sponsored search auctions, positions are naturally rival. Here, it's an option to place drugs on the same tier, and in the data, we see that a large majority of formularies exhibit this parity.

E Suggestive Evidence of History-Dependent Demand Based on Firm Behavior

As alluded to in the background section, several types of firm behavior are also suggestive of history-dependent demand. These include entry decisions surrounding me-too and incremental drugs and the amount firms pay to shorten review time by a few months. The observations also suggest that history-dependent demand may have important effects beyond pricing impacts.

¹¹³ \mathcal{F} in the notation of my empirical model.

Previous studies have shown that the entry of me-too and incremental drugs are time-dependent. Me-too drugs are branded drugs that have similar properties to existing drugs on the market, and tend to be introduced by a non-incumbent firm, whereas incremental drugs are tweaks to an existing incumbent drug introduced by the incumbent firm. For me-too drug entry, Gilchrist (2016) finds potentially causal evidence that longer patent terms on first-in-class drugs equate to more me-too entry.¹¹⁴ He argues this is due to the threat of having to compete with generics, an effect that could be amplified by any frictions in switching from a generic to a newly introduced branded drug.

Timing also appears to matter for incremental drug entry. Huckfeldt and Knittel (2011) document that they tend to be introduced and heavily advertised in the two years before patent expiration. These occur at around the same time as the price increases documented by Caves et al. (1991), consistent with the practice of “soft switches” discussed in industry circles. A soft switch is when a company introduces an incremental version of a basic drug with little patent term left and raise prices on the basic drug in hopes of establishing users on the incremental drug, which usually has longer protection remaining. A more extreme strategy is the hard switch, where a company pulls the basic drug altogether, as was the case with Namenda, an Alzheimer’s treatment. Both strategies are valuable, according to CEOs and analysts, because very few people move from the incremental drug onto the generic version of the basic drug, consistent with the findings in this paper. Recently, the FTC criticized these tactics as anti-competitive in an amicus brief on the case *Mylan vs. Warner Chilcott*.

A final piece of observational evidence is the size of first-mover advantage in the industry, evident in the sale value of FDA priority review vouchers. Ridley and Régnier (2016) discuss the sizable value of FDA priority review vouchers, which were introduced by Congress in 2007 as a reward for firms that commercialize treatments for rare diseases. Each voucher gives firms the right to convert a standard FDA review, a ten-month process mostly for me-too drugs, into a priority review, a six-month process usually reserved for drugs deemed to be major improvements. These vouchers are transferable, and five have been sold at prices ranging from \$67.5 million to \$350 million.¹¹⁵ For some context, one of the best-selling drugs of all time, Lipitor, made an estimated \$583 million in its entire first year, probably around \$65 million in its first four months, given its sales trajectory.¹¹⁶ Therefore, the amounts paid for these vouchers suggest that there must be other factors beyond just the mechanical four months worth of sales and the discounted value of extra sales from the extended patent term. One key factor could be demand-driven first-mover advantage.¹¹⁷

F Additional Graphs and Tables

¹¹⁴He uses delays between initial patenting and first clinical trial as a source of variation.

¹¹⁵See Appendix ?? for a more detailed overview of these transactions.

¹¹⁶Sales estimates reported in NY Times article <http://www.nytimes.com/1998/02/27/business/new-ranking-on-drug-sales-in-us-in-97.html>

¹¹⁷Ridley and Régnier (2016) offer a simple estimate suggesting that a third to a half of the voucher value comes from competitive considerations.

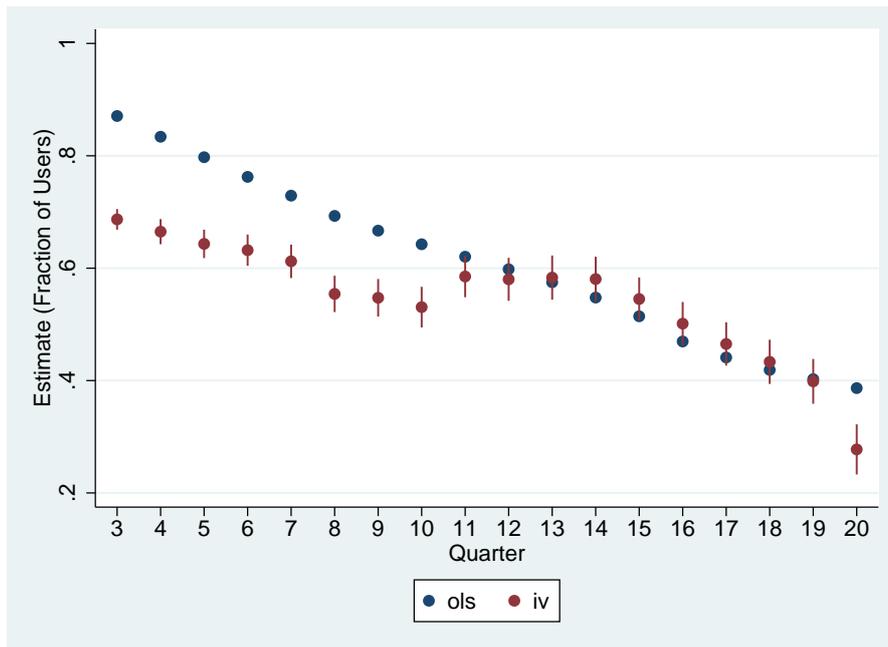
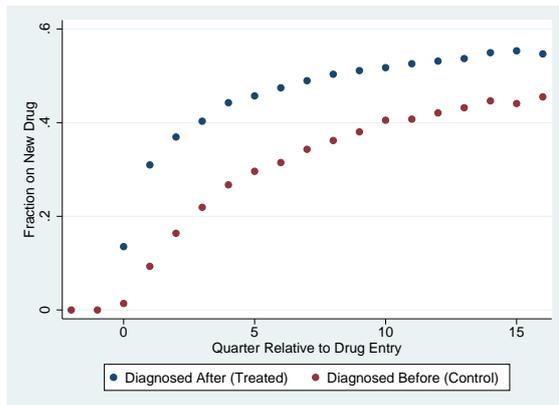
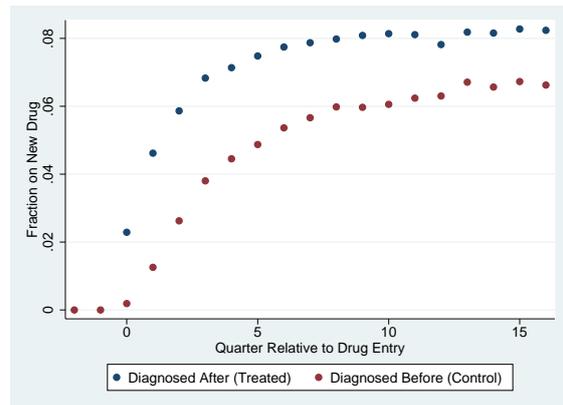


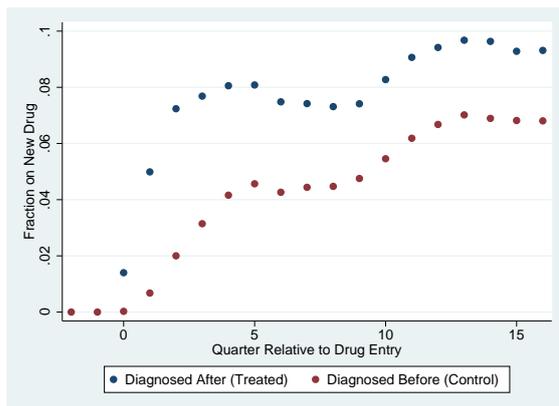
Figure 18: A plot of the β_t coefficients from Equation (1), with and without instrumenting with treatment group indicator. The data sample is the unbalanced panel of MarketScan individuals. The results suggest significant inertia effects that decrease with time, with pooling generating some non-monotonicities that are not present when analyzing individual entry events.



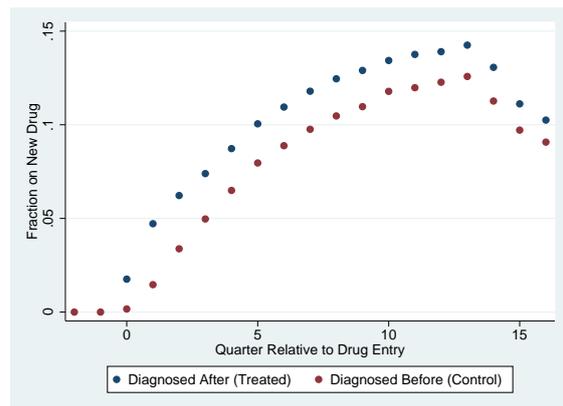
(a) Lipitor



(b) Zetia



(c) Crestor



(d) Vytorin

Figure 19: Graphs depicting differences between treatment and control groups for major entry events from 1996-2013 in the anti-cholesterol market. Treatment group contains patients who start on treatment in the 180 days after a drug launches and the control group contains patients who start in the 180 days before launch.

	t+3	t+6	t+9	t+12	t+15
Initially Chose Drug	0.815*** (0.0376)	0.786*** (0.0557)	0.628*** (0.0672)	0.769*** (0.0751)	1.057*** (0.0784)
Observations	540903	414062	320228	256722	186507
R^2	0.180	0.083	0.055	0.037	.

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 8: Estimating Equation (1), but using a 60-day window each side of an event to classify treatment and control users rather than the baseline 180-day window. Comparable to Table 2

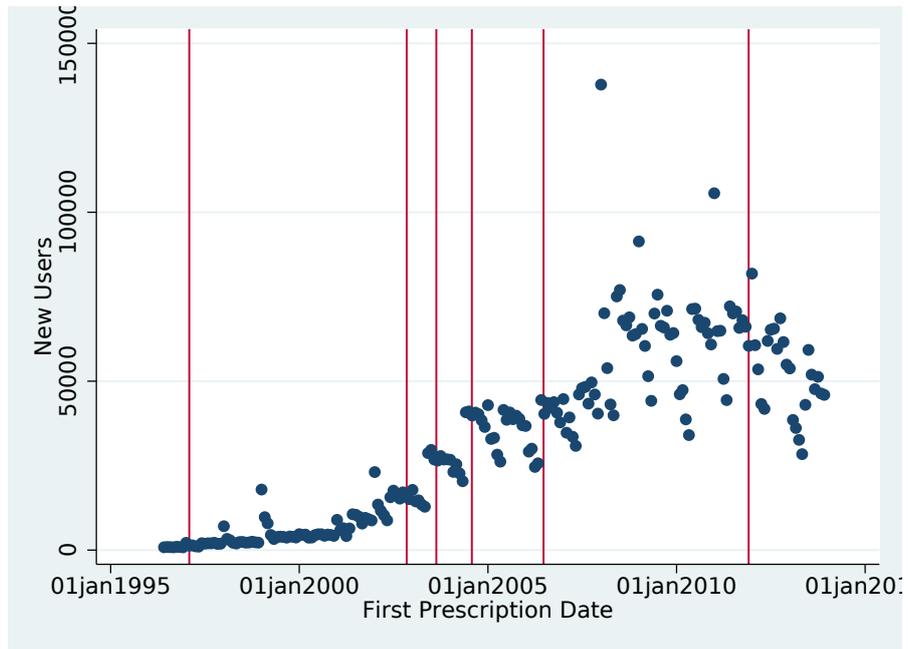
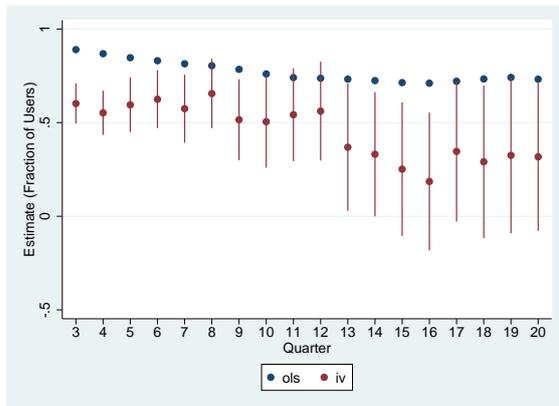
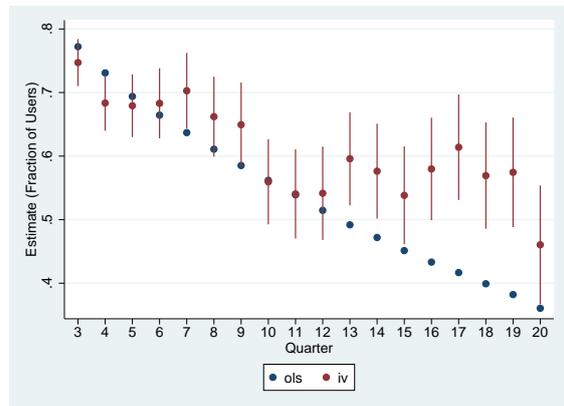


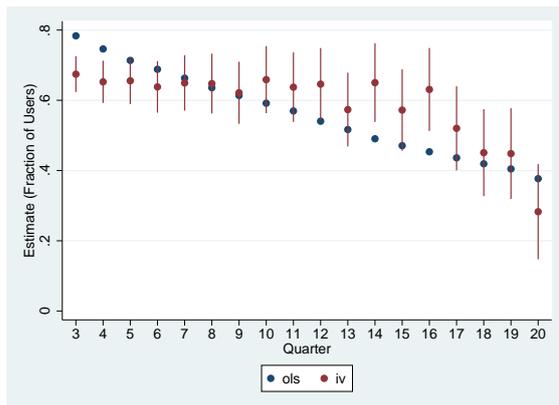
Figure 20: A plot of new patients starting on cholesterol treatment each month in the complete sample. The vertical lines highlight the entry of branded Lipitor, Crestor, Zetia, generic Zocor, and generic Lipitor. The trend generally follows the overall increase in users covered in the MarketScan data, but does not exhibit noticeable spikes around entry events.



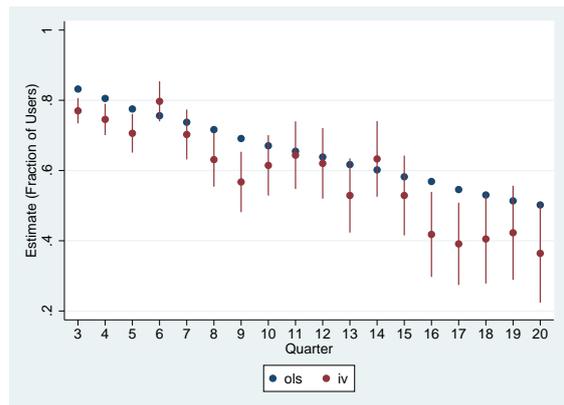
(a) Multiple Sclerosis



(b) Diabetes



(c) Hypertension



(d) COPD

Figure 21: Graphs depicting OLS and IV estimates of β_t in Equation (1), for branded drug entry events in other chronic drug markets. For multiple sclerosis, I include all new product entry events, including interferons that are introduced that have the same generic name as existing drugs on the market. The results are noisier for most of these, as there are fewer people in the dataset, but the general pattern of inertia appears to hold to varying degrees across areas.

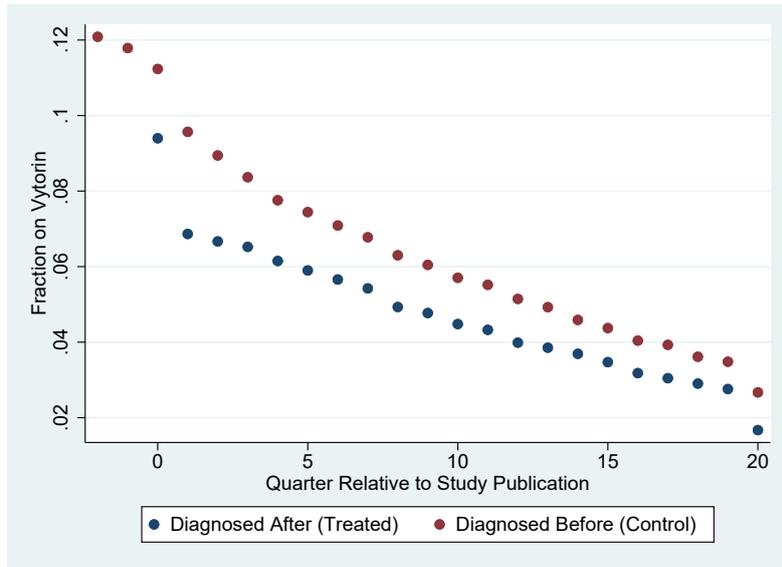


Figure 22: A plot of Vytorin market share for cohorts that start before and after the announcement of the ENHANCE trial results, which showed that Vytorin was not superior to other statins, as was thought to be the case based on LDL reduction properties. The after group (blue) has a persistently lower share of patients on Vytorin, again suggestive of history-dependence in demand. The graph shows that other types of variation in initial conditions, in this case medical information, can also have lasting impacts on choice patterns.

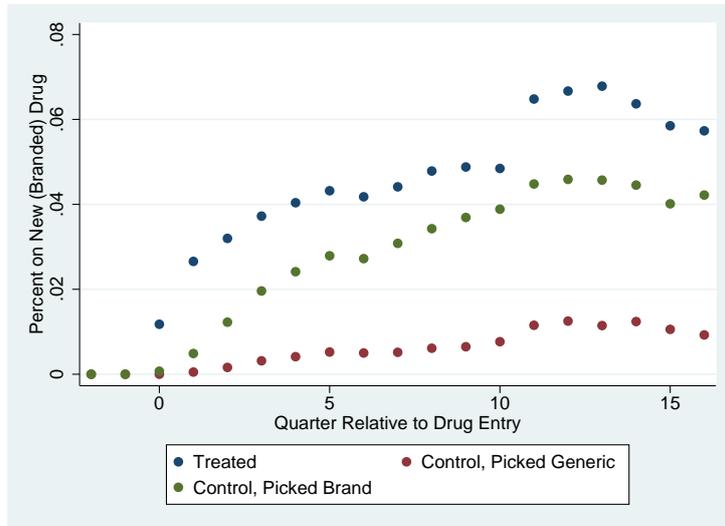


Figure 23: Entry event analysis with control group breakdown for branded drug entries. Users already on generic drugs, mostly Mevacor for branded entries in this period, are much less likely to adopt the new branded drug relative to control users who picked an existing branded drug.

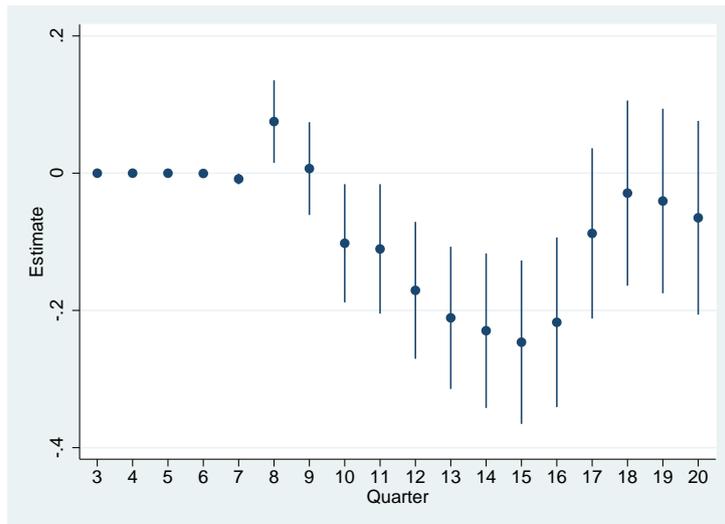


Figure 24: A plot of the estimated coefficients from Equation (1), using the Vytorin entry event but with generic Zocor usage as the outcome instead of Vytorin usage as the outcome. The graph shows that the treatment group, those pushed towards starting on Vytorin, are less likely to adopt generic Zocor, until the announcement of the ENHANCE trial around quarter 15, at which point patients on Vytorin start to switch to generic Zocor.

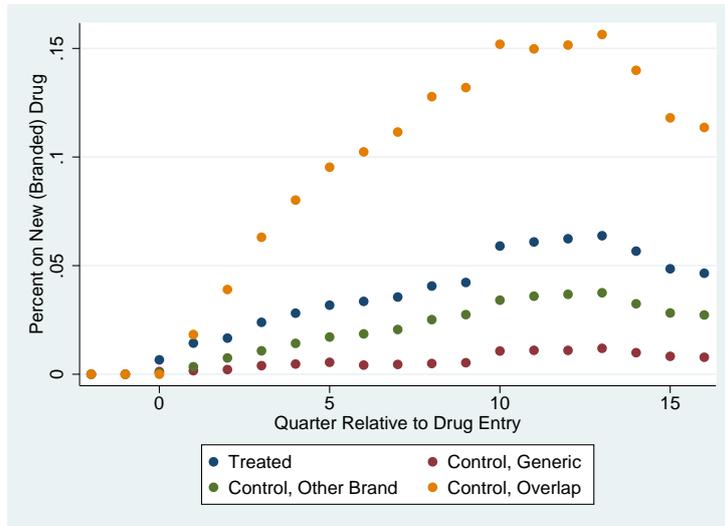


Figure 25: Entry event analysis with control group breakdown for combination therapies, including Vytorin. Patients on medication that overlap with the components of the combination are significantly more likely to adopt the combination, whereas patients already on generics rarely adopt the new drug.

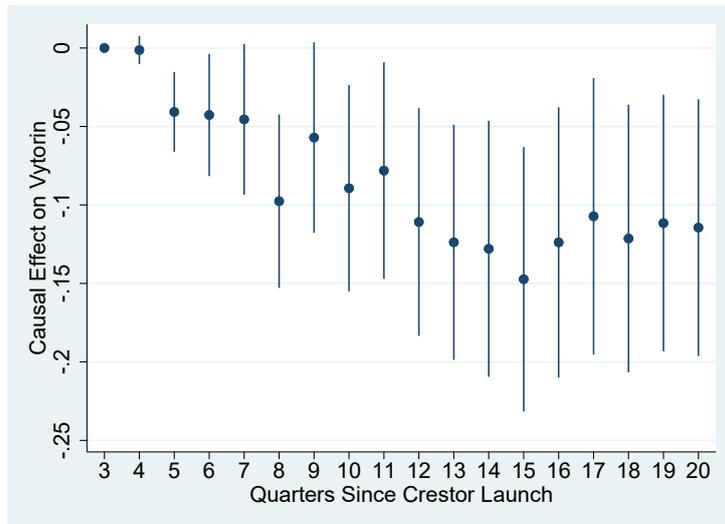


Figure 26: A plot of the estimated coefficients from Equation (1), using the Vytorin entry event but with generic Zocor usage as the outcome instead of Vytorin usage as the outcome. The graph shows that the treatment group, those pushed towards starting on Vytorin, are less likely to adopt generic Zocor, until the announcement of the ENHANCE trial around quarter 15, at which point patients on Vytorin start to switch to generic Zocor.

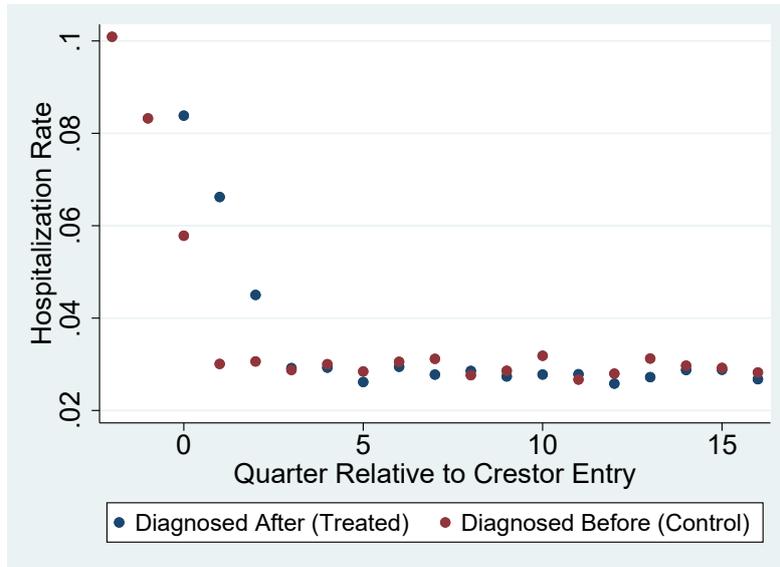
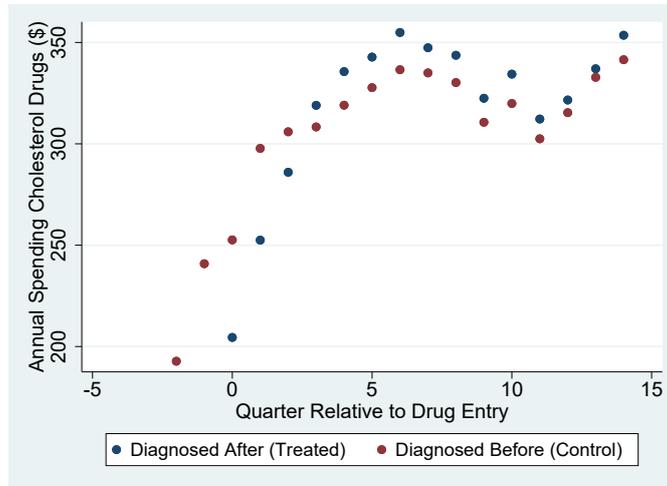
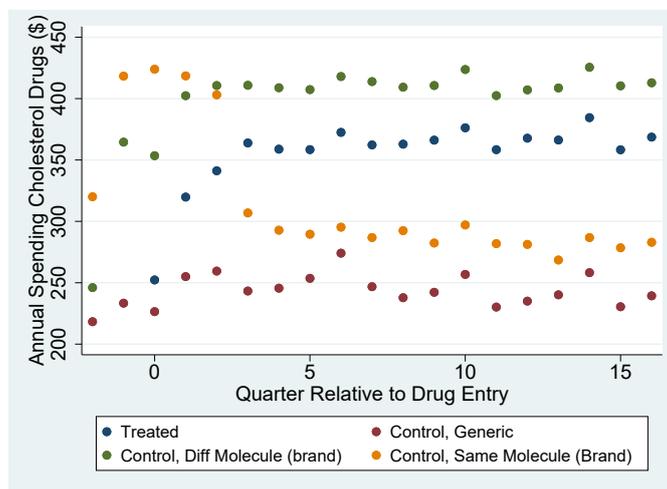


Figure 27: A plot of hospitalization rates by treatment and control group for the entry of Crestor in August 2003. The aim is to study the long-run health impacts of Crestor, which we can uncover given random initial conditions and strong inertia, as discussed in Section C.2. The two groups generally have similar long-term rates of hospitalization, suggesting minimal difference between Crestor and the existing standard of care before.



(a) Spending (Livalo Entry)



(b) Spending Breakdown (Generic Zocor)

Figure 28: Plots of average cholesterol drug spending for treatment and control groups for (a) the entry of Livalo in 2010 and (b) generic Zocor in 2006. The Livalo example illustrates initial savings from some fraction of users taking a cheaper branded drug (Livalo) versus the general baseline, but this difference reverses once generic Lipitor enters and Lipitor users, who comprise a larger share of the control group, switch to generics. Therefore, keeping users on Lipitor in the short-term may save money in the long-run. The generic Zocor breakdown illustrates a similar idea, by breaking down spending by group. Control patients who pick a different branded drug (mostly Lipitor and Crestor) start off spending less than those who start on branded Zocor, but spend much more over time relative to all other groups. Control group patients who start on branded Zocor quickly converge towards the average spending levels of the group on generic Mevacor, which was already on the market.

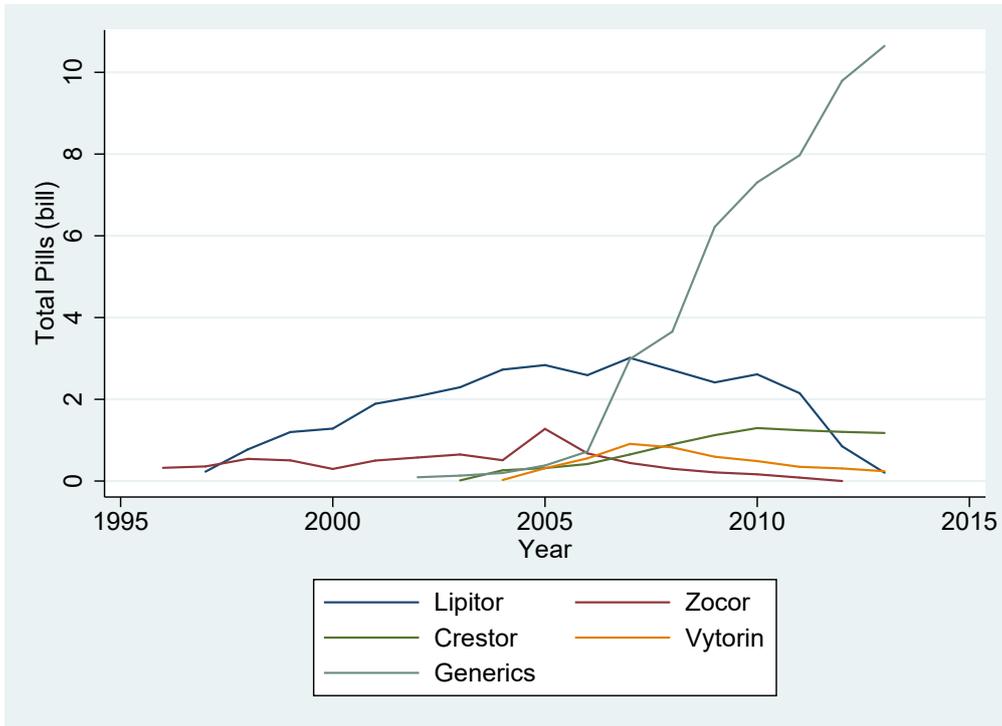


Figure 29: A raw breakdown of total pills sold for the three major statins, using estimates from MEPS. Unlike Figure 4, I aggregate over all treatment intensity levels and use a quantity-based measure. Striking patterns include that fact that Lipitor begins to dominate the market right away after its launch in 1997. Another key pattern is that generic Zocor and generic Lipitor begin to take over the market later in the period. Another aspect to note is the brief popularity of Vytorin, a combination of Zocor and Zetia, from 2006 to 2008, which is mostly used for high intensity treatment.

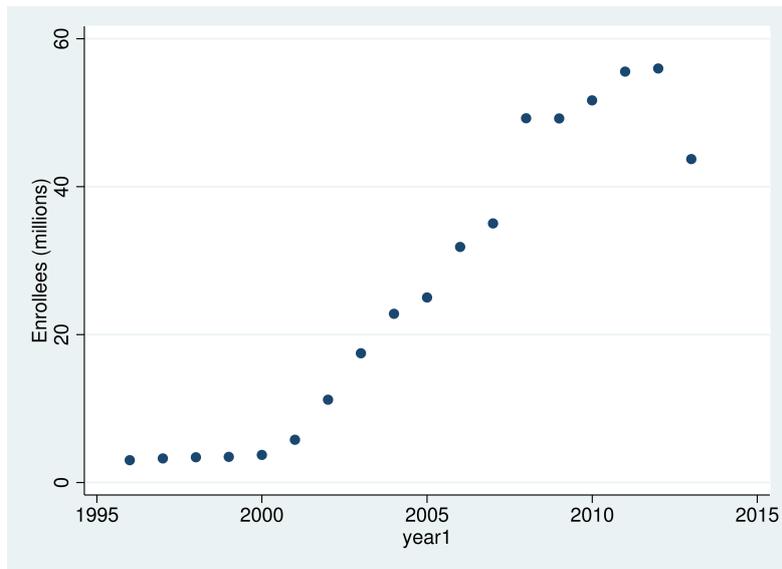


Figure 30: A plot of the number of people covered by MarketScan data over the sample period. The number increases significantly over the period, reaching a peak of over 50 million individuals in the 2008-2012 period.

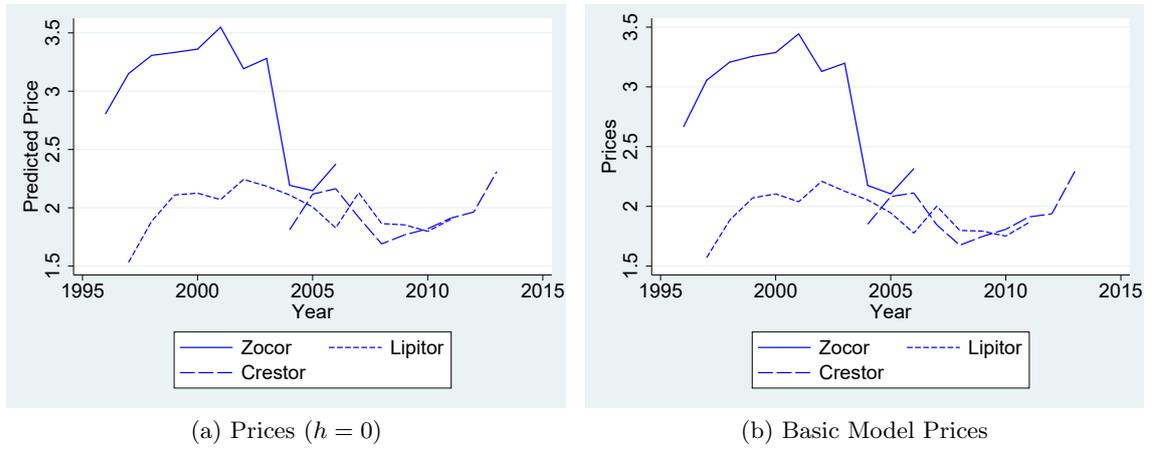


Figure 31: Counterfactual prices under a scenario with no carryover in previous market share to entering generics ($h = 0$). The two price series are almost identical, as effects from generics not benefiting from existing market share of the branded drug are minimal. This is reflective of the inherent competitiveness of generic medication and the size of inertia protecting branded competitors.