No Free Lunch? Welfare Analysis of Firms Selling Through Expert Intermediaries

Matthew Grennan*, Kyle Myers†, Ashley Swanson‡, Aaron Chatterji§

August 4, 2020

Abstract

We study how firms target and influence expert intermediaries, and the welfare impact of banning those relationships. In the case study we investigate, manufacturers of statins, a class of cholesterol-lowering drugs, provide meals and other payments to physicians. Leveraging variation in exposure to spillovers from academic medical centers' conflict-of-interest policies for identification, we estimate significant heterogeneity in the effects of payments on prescribing, with firms targeting highly responsive physicians. Payments offset the negative effects of oligopoly pricing and other frictions on utilization, but at great expense to consumers and insurers because payments promote high-price branded drugs. To understand the net effects of payments in the presence of various factors that may drive a wedge between physicians' decisions and patients' best interests, we introduce a decision error into our framework and explore the assumptions under which payments benefit consumers. We calibrate this decision error using clinical trial results on statin effectiveness for a similar population. This exercise suggests that, in the case of statins, firm payments to physicians benefit consumers due to significant under-prescribing at baseline.

*University of Pennsylvania, The Wharton School & NBER, grennan@wharton.upenn.edu

†Harvard Business School, kmyers@hbs.edu

‡Columbia Business School & NBER, ats2180@gsb.columbia.edu

§Duke University, The Fuqua School & NBER, ronnie@duke.edu

The data used in this paper were generously provided, in part, by Kyruus, Inc. We gratefully acknowledge financial support from the Wharton Dean's Research Fund, Mack Institute, and Public Policy Initiative. Gi Heung Kim and Donato Onorato provided excellent research assistance. Jorge Acosta, Abby Alpert, John Asker, Alexandre Belloni, Colleen Carey, David Dranove, Pierre Dubois, Gautam Gowrisankaran, Amit Khandelwal, Ian Larkin, Robin Lee, Aviv Nevo, Brad Shapiro, Amanda Starc, Bob Town, Heidi Williams, and numerous seminar and conference audiences provided helpful discussion and feedback. Any errors are our own. Some covariates were calculated (or derived) based on data from The Nielsen Company (US), LLC and marketing databases provided by the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. This paper subsumes an earlier working paper titled "Physician-Industry Interactions: Persuasion and Welfare."

1 Introduction

In many markets, consumers obtain expert advice before making a purchase decision. This is especially true in markets where decisions are complex or have large stakes. Firms often seek to influence those experts, and there is a growing body of empirical evidence from a variety of sectors—including insurance, financial services, and health care—that firm influence on experts' recommendations can harm consumers (Anagol et al. 2017; David et al. 2010; Egan et al. 2019; Robles-Garcia 2020). In the health care context we study, physicians receive payments from pharmaceutical and medical device manufacturers. Arguably, the interactions that accompany those payments may provide valuable information about promoted products. However, concerns about conflicts of interest have led some states and numerous academic medical centers (AMCs) to enact policies to ban or limit payments and interactions between firms and physicians (King and Bearman 2017; Larkin et al. 2017). Despite large potential financial and health stakes, little is known about the effects of such policies.

Payments from firms to physicians during "detailing" interactions have long been a key form of drug promotion. More than 85 percent of pharmaceutical marketing expenditures are targeted to physicians (Pew Charitable Trust 2013), and several studies have found a positive average association between payments and pharmaceutical prescribing (Mizik and Jacobson 2004; Datta and Dave 2016; DeJong et al. 2016; Yeh et al. 2016).[1] The policy implications of these associations are difficult to interpret in light of the well-documented facts that physician treatment behavior varies widely (e.g., Cutler et al. 2019), and that pharmaceutical firms spend large amounts of time and money targeting promotions to specific physicians (Fugh-Berman and Ahari 2007). That is, payments are not allocated randomly. Several recent papers have approached the issue of physician selection using specifications with physician fixed effects (Agha and Zeltzer 2019; Carey et al. 2020; Shapiro 2018a). This approach is instructive for estimating certain treatment effects, but it has two limitations for evaluating the impact of policies that ban or restrict payments. First, a ban entails eliminating all payments from firms to physicians, and the effect of the overall steady-state payment relationship may be much larger than the within-physician effect of an incremental payment. Second, the factors that make a given physician an attractive promotional target for a firm will generate the usual concern about selection on expected prescribing *levels*, but may also generate selection on the expected *slope* of the prescribing response to payment. If there is significant variation in how physicians respond to firm payments, this may imply large

[1] See, e.g., Spurling et al. (2010) and Kremer et al. (2008) for reviews of early research on this topic. There is also a marketing literature focused on on the causes and consequences of pharmaceutical detailing (Chintagunta and Manchanda 2004; Manchanda and Honka 2009; Narayanan and Manchanda 2009; Guo et al. 2020).

discrepancies between any single average treatment effect (e.g., a local average treatment effect) and the full distribution of effects across physicians (Heckman et al. 2006).

In the first part of this paper, we address these challenges by introducing a new instrumental variables (IV) strategy to estimate heterogenous physician responses to firm payments. Our IV strategy leverages differential exposure of individual physicians to spillovers from AMC policies to generate quasi-exogenous variation in who firms target with payments. We estimate large variation in treatment effects across physicians, and find that pharmaceutical firms target physicians based on this variation.

The second part of this paper adds a combination of structural modeling and calibration from external data sources (in particular, evidence from clinical trials) to move from our estimated treatment effects toward a fuller understanding of the equilibrium welfare effects of payments. As noted by recent theoretical work on conflicts of interest in health care and finance (Inderst and Ottaviani 2012), firm interactions with experts tend to occur in markets with other frictions. Here, understanding how treatment effects map into equilibrium outcomes requires understanding how they interact with oligopolistic competition, insurance, and any other frictions—such as imperfect agency or behavioral biases (Baicker et al. 2015; Currie and MacLeod 2020)—that might drive a wedge between the treatment a physician chooses and the treatment that maximizes patient welfare.

We illustrate our research design and modeling strategy using a case study in the market for statins, an important class of cardiovascular drugs, from 2011–2012. We combine prescribing data from the Medicare Part D prescription drug insurance program with payment data at the physician-drug-year level. We also merge numerous other physician, hospital, and market characteristics. The two promoted (branded) statins in our sample (Pfizer's Lipitor and AstraZeneca's Crestor) had prices around seven times those of generic alternatives, yet still made up nearly 40 percent of statin prescribing in 2011. Notably, these particular drugs were the first "strong statins" on the market, and were shown to generate larger reductions in cholesterol for some patients. They were also heavily promoted drugs: over 75 percent of prescriptions in our data were written by a physician who received a meal from at least one of their manufacturers (meals and related interactions represent almost 98 percent of payment instances between pharmaceutical firms and physicians). Thus statins are interesting in and of themselves, as one of the largest-selling drug categories in history, but they also provide a representative example of a market where firms may use payments to influence utilization in favor of their more expensive products, which might offer little value (to consumers) relative to cheaper substitutes (Scott Morton and Kyle 2012).

We estimate the causal effects of firm payments on physician prescribing by leveraging variation in physicians' exposure to AMCs' conflict of interest policies, which restrict firms'

ability to provide payments to affiliated physicians. Motivated by geographic economies of scale in firms' marketing efforts, we document that the effects of these policies spill over to other physicians who are unaffiliated with the AMC but happen to practice nearby.[2] The pattern of how these policy spillovers impact payments matches our motivating theory: spillovers are stronger when a larger proportion of a region's cardiologists are affiliated with the AMC, and they are weaker for physicians who are located farther from the AMC. We also document that, for subsets of physicians where we expect (and observe) these policy spillovers to have no "first stage" effect on meals, the spillovers also have no "reduced form" effect on prescribing, providing more confidence that our instruments for meal receipt are not correlated with unobservable determinants of prescribing (Bound and Jaeger 2000; Altonji et al. 2005; Angrist et al. 2010).

We use this quasi-exogenous variation in payments in a semi-parametric instrumental variables procedure to estimate the full distribution of marginal treatment effects (MTEs; Heckman and Vytlacil 2007). In order to gain the predictive power needed to precisely estimate MTEs, and also control for physician- and market-level differences, we include flexible functions of a large number of potentially relevant control variables in our estimation. This creates a dimensionality and sparsity problem, which we address by drawing on the recent literature at the intersection of machine learning and econometrics (Belloni et al. 2017; Chernozhukov et al. 2018). In particular, we use Lasso regressions to select the most powerful set of predictors, while employing sample splitting to ensure that our estimates are robust to errors in the variable selection process. To control for substitution between competing drugs, we embed the MTE estimation routine in a discrete choice random utility model of demand for statins as a function of drug quality, payments, and patients' out-of-pocket prices. We identify substitution patterns and price sensitivity by leveraging the expiry of Lipitor's patent at the end of 2011 and the ensuing generic entry as an exogenous shock to both the choice sets physicians faced and to consumers' out-of-pocket prices.

This analysis yields an important new result: there is dramatic variation in physicians' responsiveness to payments. Our estimates imply that a meal-based detailing relationship increases promoted statin prescribing by 45 percent for the average physician, which is roughly equivalent to the impact of a $40 price decline or half of a standard deviation in the prescribing heterogeneity across physicians. However, for a physician in the 90th percentile, the effect is equivalent to one standard deviation, while the effect in the 10th percentile is not statistically different from zero. We also find that firms target physicians who: (i) have larger

[2]This "spillovers" identification strategy is similar to that in Hastings et al. (2017), which relies on variation in sales force exposure driven by the characteristics of other nearby investors. See also Waldfogel (2007) for a broader discussion. Larkin et al. (2017) focuses on the direct effects of these policies and estimates significant reductions in prescribing of detailed drugs at the institutions that impose them.

expected treatment effects, (ii) would otherwise prescribe below-average amounts of the firms' drugs, and (iii) have larger patient panels. As a result, among physicians who receive meals, the incremental revenue due to meals is large. For most physicians not targeted, though, we cannot reject the null that a meal would lead to zero extra revenue.

Translating these treatment effects on prescribing into welfare effects requires not only the substitution patterns estimated in our demand model, but also a model of how prices are set in equilibrium and a way to map prescribing outcomes into consumer welfare. To address the former, we estimate a model of price negotiation between upstream manufacturers/distributors and insurers to capture the forces driving the point-of-sale prices that insurers pay for pharmaceuticals. Our bargaining parameter estimates are consistent with branded firms receiving a large portion of the surplus they create, while competition among many firms drives down margins on generics dramatically.

Deriving consumer welfare implications from changes in prescribing is difficult because even in the absence of payments and pricing frictions, prescribing could be suboptimal due to variation in physician information and skill (e.g., Currie and MacLeod 2020), imperfect agency not driven by detailing/meals (Jacobson et al. 2006), or various behavioral biases (Baicker et al. 2015). Payments could reinforce or counteract such frictions.[3] This is closely related to the question, endemic to the marketing literature, of whether advertising is "informative" vs. "persuasive." To address this, we introduce a "decision error" parameter that may be positive or negative and that drives a wedge between decision utility and realized welfare (similar to the error in Baicker et al. (2015), but intended to capture non-behavioral errors as well). We simulate the welfare impact of a payment ban for a wide range of decision error values, and we combine our revealed preference utility estimates with clinical data on statin effectiveness to calibrate the sign and magnitude of the decision error our case study.

Viewing the effects of physician-industry relationships through the lens of the structural model yields several additional insights. First, simulating a counterfactual meal ban shows that the equilibrium effect of meals on prescribing is to increase statin use by around five percent and use of the focal branded statins by 29 percent on average. These numbers are smaller than the above estimated treatment effect of a meal on prescribing because not all physicians receive meals (in part because many have small expected responsiveness), and because many physicians who do receive meals get them from both competing firms so that business stealing effects cancel each other out. Second, high branded statin prices lead to prescribing below the efficient level in a world without meals. Our estimated model suggests

[3]In the former case, payments represent harmful kickbacks. For example, Novartis recently paid nearly $700 million to settle a whistleblower suit regarding physician payments under federal anti-kickback law (Morgenson 2020). In the latter case, payments are helpful, but expensive, nudges. Statins are still often cited as a class of drugs that is underprescribed relative to clinical guidelines (Walter 2020).

that payments increase prescribing to near the efficient level (according to revealed preference demand estimates), though at high cost to consumers and payers.

Turning to results that depend upon baseline under-/over-prescribing: we find that if decision utility reflects true consumer utility, then meals result in large surplus gains to producers, negatively impact consumers, and have a small negative impact on total surplus on net. If, however, behavioral or other frictions bias consumers' revealed willingness to pay for statins downward by a substantial amount, then consumer surplus increases in the presence of meals. In this way, our framework yields a threshold value, above which meals are welfare-reducing (and vice versa). In Section 5.5.1, we find that this threshold is well above the decision error value that would equate average willingness-to-pay in our framework with conservative estimates of the life-year gains from statin regimens among indicated patients in the clinical literature, and we discuss the assumptions necessary to apply these outside estimates to our empirical setting. When we calibrate the model to the clinical data in this way, the welfare impact of a meal payment ban is substantial – the total surplus effect is similar in magnitude to that of introducing generic atorvastatin (one of the largest generic introductions in history).

In addition to detailed empirical estimates for the important case of statins, this paper contributes broadly useful new instrumental variables, and a framework for estimating heterogeneous treatment effects of pharmaceutical firm payments to physicians and mapping those treatment effects into equilibrium welfare effects. Our findings add to the literature on potential conflicts of interest among expert intermediaries across a range of markets (Anagol et al. 2017; Egan et al. 2019; Levitt and Syverson 2008; Robles-Garcia 2020; Schneider 2012) and in particular the literature on drivers of physician treatment recommendations (Johnson and Rehavi 2016; Gruber and Owings 1996; Iizuka 2012; Clemens and Gottlieb 2014; Ho and Pakes 2014). Our focus on heterogeneity in treatment effects of targeted promotional activity adds new elements to a small but growing set of studies that consider the equilibrium effects of expert inducements in imperfectly competitive markets (e.g., Egan et al. 2020; Robles-Garcia 2020), and of firm advertising in markets with health consequences (see, e.g., Dubois et al. (2018) regarding junk food advertising, and Shapiro (2018b) and Sinkinson and Starc (2019) regarding direct-to-consumer advertising of drugs). Finally, our approach to mapping demand into welfare in the presence of unobserved decision frictions offers a new path forward in cases such as ours where outside data on the benefits of a product are available. This approach adds to a growing literature that has so far required unique data on welfare-relevant endpoints (Abaluck et al. 2020) or on which decision-makers are unlikely to be biased by such frictions (Allcott and Taubinsky 2015; Bronnenberg et al. 2015; Handel and Kolstad 2015; Handel and Schwartzstein 2018).

The remainder of the paper is as follows: Section 2 describes our empirical setting and summarizes the high-level patterns in our data in order to motivate our empirical model. Section 3 illustrates the intuition underlying our causal identification strategy. Section 4 steps through our demand model, estimation approach (including MTE and machine learning methods), and results. Section 5 presents our welfare framework, supply-side model, and the results of counterfactual simulations of a ban on physician-firm payments. Finally, our concluding Section 6 discusses the extent to which one might draw cautious policy implications from our estimates and extend the data and approach in future research to better inform policy.

2 Setting, Data, and Summary Statistics

This Section first describes institutional details of the market we focus on (the Medicare statin market in 2011–2012). It then describes our sources for data on (1) drug quantities and prices, (2) payments from firms to physicians, and (3) other regional characteristics. Finally, it outlines our sample restrictions and summary statistics.

2.1 Setting: Medicare Statin Market, 2011–2012

With prescription drugs accounting for more than 15 percent of personal health care expenditures, and with 72 percent of that attributed to branded drugs, the potential financial and health consequences of branded drug manufacturers' payments to physicians are significant (ASPE 2016; Kesselheim et al. 2016). In this study, we focus on cardiologists' prescriptions of statins in the Medicare Part D program for the elderly in the U.S. in 2011 and 2012. This sample and time horizon are useful for several reasons: (1) We observe all payments for both branded statin manufacturers in these years. Pfizer (which produces Lipitor) and AstraZeneca (which produces Crestor) accounted for 49 percent and 33 percent of statin revenue in Medicare Part D in our sample in 2011, respectively. (2) Statins are an important class of drug in their own right. While Lipitor was on patent, it was the best-selling drug in the history of pharmaceuticals. (3) These statins were each the chief source of revenue from cardiologists' prescribing for these two firms, with Lipitor accounting for 84 percent of Pfizer's cardiologist-driven revenues and Crestor similarly accounting for 80 percent of AstraZeneca's cardiologist-driven revenues. Thus, if a Pfizer or AstraZeneca representative were taking a cardiologist out to lunch in this time period, it is very likely that statins were the focus of any drug-related discussions.[4] (4) Finally, Lipitor's patent expiration offers a

[4]Cardiologists account for 10 percent of Part D statin claims. Although cardiologists write relatively few prescriptions, they are targeted because specialist prescriptions are often sustained by primary care

large and visible shock to statin prices and substitutes, helping to identify other features of demand curves separately from payment effects.[5]

Statin medications reduce blood levels of low-density lipoprotein cholesterol (LDL, or "bad" cholesterol), and in turn reduce the risk of coronary heart disease and heart attacks. Statins are generally considered to be effective; the American College of Cardiology (ACC)'s 2013 guidelines recommended statin therapy for adults with elevated risk of atherosclerotic cardiovascular disease. Full adoption under these guidelines would have increased statin use by 24 percent (American College of Cardiology 2017). Statins are close substitutes for most patients, but atorvastatin (Lipitor) and rosuvastatin (Crestor) are available as high-intensity "strong" statins appropriate for some patients with elevated risk (ConsumerReports 2014).

The structure of Medicare Part D (see Appendix A.1 for detail on the program) implies that enrollees should be sensitive to price variation across and within branded and generic drugs. This sensitivity may be muted by various frictions, including enrollees' limited understanding of coverage and physicians' imperfect agency.[6] Part D plan issuers' strategies and profits are heavily regulated by the Centers for Medicare and Medicaid Services (CMS), but they have both motive and opportunity to constrain costs through formulary design (i.e., drugs' placement on tiers) and negotiations with drug manufacturers (Duggan and Scott Morton 2010).

2.2 Data: Physician and Drug Characteristics, Quantities, and Prices

We obtain data on physician specialties, affiliations, and demographics from the CMS Physician Compare database, which contains all physicians treating Medicare patients.[7] Each physician's practice location is matched to his or her relevant Hospital Service Area (HSA) and Hospital Referral Region (HRR) according to the Dartmouth Atlas.[8]

Prescribing data are from the publicly-available CMS Part D claims files for 2011 and

physicians, giving them an outsized impact on total prescribing (Fugh-Berman and Ahari 2007).

[5]The generic version of Lipitor (atorvastatin) became available in December 2011. The entry of this generic drug created the customary shocks to absolute and relative prices that follow the loss of exclusivity, and at a very large scale: the total Part D expenditures associated with Lipitor dropped by more than 75 percent, from $2.5 billion (13 million claims) in 2011 to $591 million (2.8 million claims) in 2012.

[6]For example, enrollees are more responsive to current prices than marginal prices, and respond disproportionately to salient coverage changes such as copay changes for entire drug classes (Abaluck et al. 2018). See Goldman et al. (2007) and Chandra et al. (2010) for helpful reviews of the literature.

[7]See: https://data.medicare.gov/data/physician-compare.

[8]See: www.dartmouthatlas.org for more. HRRs represent regional health care markets for tertiary medical care and are based on proximity to cities where major cardiovascular and neurosurgery procedures are performed. HSAs are local health care markets for hospital care and are based on aggregations of ZIP codes where individuals tend to use the same hospital. There are 3,436 HSAs and 306 HRRs in the U.S.

2012.[9] These claims data describe total prescriptions (in 30-day supplies) and spending for each prescriber-drug-year. The prescriber information includes physicians' National Provider Identifiers (NPIs), which allows us to link claims data to other data sources. Drugs are defined by brand and molecule name (if the drug is "generic," these two are equivalent). Prescriptions may vary in terms of unobserved drug dosages and formulation. However, we are unaware of any evidence that industry payments target particular dosages or formulations, so we follow prior studies in analyzing days supplied as the unit of quantity (Starc and Swanson 2020).

Using the name of the drug, we also match branded drugs in the prescribing data to their respective manufacturers using the FDA's Orange Book and match all drugs to their WHO Anatomical Therapeutic Classification (ATC) codes. The ATC codes provide a hierarchy of drug categories that reflect similarities in drug mechanism and disease intended to treat. In that way, it usefully mimics the choice sets faced by physicians. We focus on the focal drug's share of all cardiovascular (ATC code = "C") and statin (ATC code = "10AAC") prescribing within physician-year. We think of cardiovascular prescribing for a physician-year as a proxy for the total number of patients seen by the cardiologist who might potentially need a statin.

Starting with the full sample of cardiologists in the Medicare Physician Compare database, as identified by their self-reported primary specialty, we restrict our sample to "active" Medicare prescribers with at least 500 Part D cardiovascular prescriptions on average in 2011 and 2012. This is approximately the 10^{th} percentile of prescriptions per physician-year. We then restrict the sample to cardiologist-statin molecule pairs that have at least two non-zero observations (which is required to estimate the mean utility parameter). The final sample used in our analyses contains about 13,000 cardiologists. In terms of focal drugs, we focus on the six most popular statins (two branded, four generic), representing over 99 percent of Part D statin prescriptions and expenditures during 2011–2012. Appendix Table A2 details the impact of these sample restrictions on key summary statistics, which is small.

Table 1 summarizes the claim quantities and drug prices for our sample. On average, a physician in our sample writes about 3,500 Medicare prescriptions in the cardiovascular class per year, and more than 600 of these are for statins. There is significant variation across physicians that we explore in more detail below, but the averages themselves clarify several important facts about the statin market during our study. The effect of entry by generic atorvastatin in December 2011 is clear—in its first full year of availability, this new alternative accounted for roughly 27 percent of cardiologists' statin claims, while Lipitor's share was only about 8 percent.

The remaining columns of Table 1 summarize prices. Our price variables are the plan

[9] See: www.cms.gov/Research-Statistics-Data-and-Systems/Research-Statistics-Data-and-Systems.html.

Table 1: Prescribing Summary Statistics

	Prescription Count, mean		Out-of-Pocket Price ($), mean		Point-of-Sale Price ($), mean	
	2011	2012	2011	2012	2011	2012
All Cardiovascular	3,438	3,887				
All Statins	626	733				
Crestor	109	120	31.86	31.85	137.09	160.33
Lipitor	148	57	32.04	62.62	139.48	163.92
Atorvastatin		195		9.67		32.45
Other Generics (3)	399	407	4.55	3.84	13.31	10.30

Notes: Based on 123,809 cardiologist-molecule-brand-year (*djbt*) observations. Prescriptions (30-day equivalent) and prices from the Medicare Part D public use files. Out-of-pocket prices are enrollment-weighted averages of Part D enrollee cost-sharing per 30-day supply. Point-of-sale (POS) prices shown are those paid (by plans and beneficiaries) when Medicare Part D enrollees fill prescriptions. POS prices do not account for manufacturer rebates or markups paid to intermediaries (e.g., pharmacies), which are taken into account later in our modeling/estimation. See Appendices A.2 and B for details on sample and variable construction.

enrollment-weighted average point-of-sale and out-of-pocket prices per one-month supply for each drug-year from the Medicare Part D Public Use Files (one month is the modal supply per claim). In 2011, Lipitor and Crestor out-of-pocket (OOP) prices—the prices paid by the patient when filling a prescription—were about seven times those of generics. The full point-of-sale (POS) prices paid by insurers to the upstream supply chain were three to four times OOP prices, and similarly an order of magnitude higher for the branded statins relative to the generics. As in most studies of pharmaceuticals, it is impossible for us to observe confidential rebates negotiated between statin manufacturers and Part D plans, or to observe the unit price ultimately obtained by manufacturers (i.e., excluding markups applied by other supply chain intermediaries). However, average rebate data reported to CMS, taken together with several recent papers that infer average rebates and supply chain markups using SEC filings (e.g., Kakani et al. 2020; Sood et al. 2017; Yu et al. 2018; see Appendix E for details), suggest that 55–68 percent of POS prices would flow through to branded manufacturers. We incorporate rebates and perform robustness to a range of reasonable assumptions in our welfare analysis in Section 5.

In 2012, generic atorvastatin was introduced by two manufacturers with 180 days of generic exclusivity (see Appendix A.2 for details on the entry environment). Atorvastatin had significantly lower OOP and POS prices than Lipitor, but prices were still higher than those other generics due to initially limited generic competition. While other generic drugs' prices were somewhat lower in 2012 than in 2011, both Pfizer and AstraZeneca increased their POS prices in 2012. Crestor's OOP price was approximately the same in 2011 and 2012,

but Lipitor's OOP price nearly doubled as insurers removed Lipitor from their formularies thereby increasing patient cost sharing. Branded manufacturers are not passive when their drugs lose exclusivity. For example, there is evidence that Pfizer aggressively promoted a copay coupon program around this time (Aitken et al. 2018), and offered larger rebates to insurers after generic atorvastatin entry (Arcidiacono et al. 2013). Copay coupons cannot be used by Medicare Part D enrollees, so we omit them from our analysis. In our supply side estimation in Section 5.3.1, we allow for higher rebates by Pfizer in 2012 and test robustness of our results to this choice.

2.3 Data: Pharmaceutical Manufacturer Payments to Physicians

More than 85 percent of pharmaceutical marketing expenditures are targeted to physicians (Pew Charitable Trust 2013). Typically, firms provide physicians with meals and other payments as part of a "detailing" relationship. These in-kind payments and their associated interactions may allow firms to inform physicians about a drug's characteristics. They may also distort utilization in favor of a firm's expensive branded drug, which might offer little clinical benefit relative to cheaper substitutes (Scott Morton and Kyle 2012).

Although federally mandated reporting of pharmaceutical manufacturer payments to physicians did not begin until 2013, nationwide interest had been growing for some time. By 2010, several states had begun to institute their own payment limitations and/or public reporting rules;[10] a number of high-profile lawsuits found conflicts of interest between physicians and manufacturers to be a punishable offense;[11] and calls from politicians and patient advocacy groups were gaining momentum (Grassley 2009). Amidst this growing concern, a number of firms, including Pfizer and AstraZeneca, began to publicly release data on payments to physicians, often due to legal settlements (Ornstein and Grochowski Jones 2015).[12] These documents are the basis of our payments data, which were generously shared by Kyruus, Inc.[13]

[10]The District of Columbia, Maine, and West Virginia required disclosure of payments and gifts to physicians prior to our time horizon; Massachusetts, Minnesota, and Vermont required disclosure and had certain statutory gift bans (King and Bearman 2017). The Physician Payment Sunshine Act mandated disclosure nationwide at OpenPayments.CMS.gov beginning in August 2013.

[11]For example, in 2009 Eli Lilly paid a $1.4 billion fine following allegations of the off-label promotion of its drug Zyprexa (United States Department of Justice 2009).

[12]The existence of some voluntary disclosures is not entirely surprising. In 2009, the industry trade association PhRMA introduced a voluntary Code on Interactions with Healthcare Professionals limiting informational presentations to the workplace and entertainment to "modest meals," and prohibiting trips to resorts, sponsored recreation, and gifts to the physicians.

[13]The raw disclosures were published in a wide variety of formats both across firms and within firms over time. In order to account for irregularities in formatting—primarily of names—a machine learning algorithm was developed by Kyruus to create a disambiguated physician-level dataset of payments from Pfizer and AstraZeneca in 2011 and 2012. Appendix B.3 compares this data to that made publicly available

Table 2 summarizes our data on payments from firms to physicians. As shown in Panel A, payments most often involve the provision of meals. In fact, meals account for 98 percent of the payments we observe in our data. Panel B shows how the distribution of meal payments very closely maps the distribution of overall payments as well. The only exception to this is the roughly one percent of physicians receiving very large payments due to consulting, speaking, and travel fees or research grants. While this one percent is in an interesting group, we focus our analysis on meals since they are clearly the dominant form of payment in this setting.

Table 2: Payments Summary Statistics

Panel (a): Fraction of Cardiologists Receiving Payments, by Type

	Any Kind		Meal		Travel/ Speak/Consult		Research	
	Raw	Claim wgt.	Raw	Claim wgt.	Raw	Claim wgt.	Raw	Claim wgt.
Crestor	0.615	0.770	0.607	0.761	0.014	0.027	0.001	0.001
Lipitor	0.338	0.443	0.316	0.417	0.014	0.027	0.001	0.001
Either	0.685	0.782	0.669	0.766	0.027	0.042	0.002	0.002

Panel (b): Payment amount ($) if >0, by Type

		Mean	$p10$	$p50$	$p90$	$p99$
Crestor	Any	432.6	15.0	58.5	176.0	10,914.5
	Meal	81.3	15.0	54.0	159.0	540.0
Lipitor	Any	323.5	11.0	33.0	143.0	6,020.5
	Meal	51.2	11.0	25.0	120.0	313.5
Either (+)	Any	548.1	15.0	74.0	243.5	13,350.5
	Meal	97.9	15.0	64.8	203.5	596.5

Notes: Based on 25,323 cardiologist-drug obs.; prescription-weighting using 2011 claims. In Panel (a), the "Either" category reports whether the cardiologist received payments from either firms. In Panel (b), which reports the distribution of non-zero payments, the "Either (+)" category reports the sum of payments from both firms.

Sixty-seven percent of physicians, representing 77 percent of cardiovascular prescriptions, received a meal from at least one of the branded statin manufacturers. Detailing relationships are highly persistent over time: for the firm-years in our estimation sample, 73 percent of physicians receiving a meal in year t also receive a meal in year $t+1$. Further, there is not a large amount of variation in the dollar amount of meals when outliers are excluded: for 90 percent of physicians, the total dollar value of meals received from a given firm in a given year (conditional on receiving any) is less than $134. While these dollar values may seem trivial, they represent only a fraction of the total cost of the overall detailing relationship

post-Sunshine Act and finds no evidence of any major biases or censoring in our data.

(see Liu et al. 2020 and further discussion below), and research has shown that even small promotional efforts can have large effects on perceptions of drug quality (Grande et al. 2009).

Motivated by these patterns in the data and institutional details, we focus most of our analysis on an indicator for whether a physician ever received a meal from a manufacturer in our data, as our measure of a physician's payment from that manufacturer. During these meals (and other interactions for which meals may proxy), sales representatives target prescribers with product presentations regarding safety, efficacy, side effects, convenience, compliance, and reimbursement. These in-the-field sales representatives are considered "the most expensive and, by consensus, highest-impact promotional weapon" in pharmaceutical firms' arsenals (Campbell 2008). The cross-sectional indicator for a meal seems to comport best with our goal of estimating the treatment effect of *any* relationship to inform welfare simulations of a ban on *all* such relationships. In Appendix G.5, we show that our results are robust to alternative definitions of the payment relationship.[14]

2.4 Variation in Prescribing and Payments

Figure 1 documents the geographic variation in utilization and meal payments across the U.S. Aggregating to the HRR level, Panel (a) plots the utilization of strong statins, and Panel (b) plots the share of cardiologists that receive meals from either of the firms in our sample. Both show significant heterogeneity. For prescribing, the 90th percentile HRR is about six times more likely than the 10th percentile HRR to prescribe a branded strong statin. For payments, the 10th percentile HRR has about 40 percent of prescriptions written by a physician who received a meal from at least one of the branded statin firms, while for the 90th percentile more than 95 percent of prescriptions come from a physician who received a meal.

Despite the large geographic variation in both prescribing and payments, no strong visual pattern emerges in how the two may be correlated. This is borne out in the table at the bottom of Figure 1, which shows the distribution of the share of cardiovascular prescriptions written for Lipitor and Crestor, split by whether the physician received a payment from the focal firm. The two distributions are nearly identical. If anything, there is slightly more prescribing of the focal drug among physicians who do not receive payments from its manufacturer.

[14] Briefly, we find no meaningful differences in treatment effects as a function of meal dollar value, and our results are similar if we flag meals based only on what we observe in 2011, and if we instead use an indicator for receipt of *any* type of payment (e.g., meals, consulting, speaking, travel, or research).

Figure 1: Variation in Prescribing and Payments in 2011

(a) Cardiov. share (p.p.)

(b) Share of docs with Meals (p.p.)

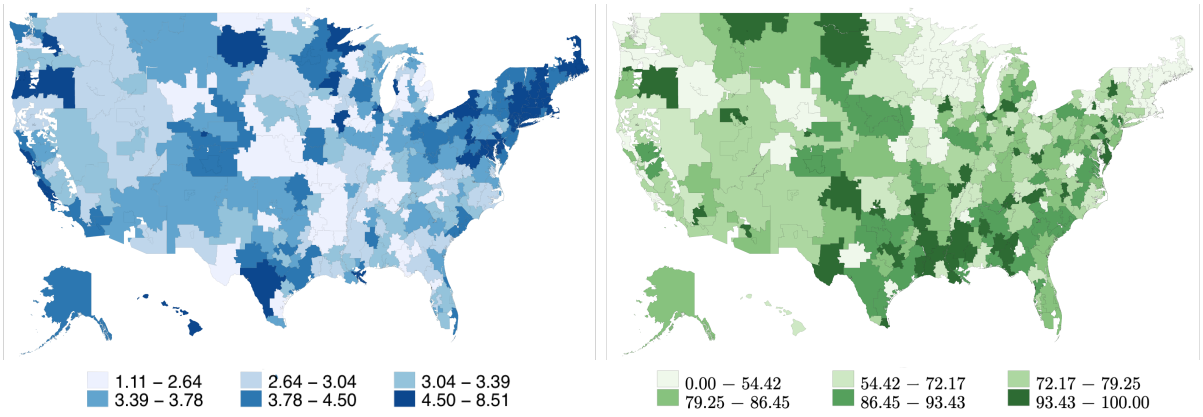

(c) Cardiov. share distribution (percentage points), by Meal status

		mean	s.d.	$p10$	$p50$	$p90$
Cardiov. share (p.p.), Crestor or Lipitor	$1_{\$>0}$, same firm	3.70	2.58	1.17	3.11	6.96
	$1_{\$=0}$, same firm	3.98	2.72	1.29	3.32	7.54

Notes: (a) the 2011 HRR-level averages of doctor-level cardiovascular shares for Crestor plus Lipitor. (b) the HRR-level share of doctors receiving meals from either AstraZeneca or Pfizer. (c) the 2011 distribution of doctor-level cardiovascular shares for Crestor or Lipitor, split by whether the "same firm" that produces the drug gave the doctor a meal. All numbers are in percentage points.

2.5 Data: Other Regional Characteristics

Our empirical strategy to identifying the effect of these meals will rely on variation induced by policies that affect some regions more than others. To help ensure that this cross-sectional variation in physician-firm payments is due to these policies and not underlying variation in other factors (e.g., preferences, market structures), and to provide statistical power, we will control for a rich set of observable physician-level and regional characteristics.

We supplement the cardiologist-specific features we can observe with: (1) ZIP code-level measures of local TV advertising for each of the two branded drugs from the Nielsen AdIntel database; (2) Hospital Service Area- (HSA) and Hospital Referral Region- (HRR) level aggregations of physician counts, practice counts, and utilization measures (e.g., total claims, cardiovascular claims per physician) from the Physician Compare and Part D data; (3) HSA- and HRR-level Medicare Advantage eligibility and penetration from CMS data; (4) HSA- and HRR-level measures of uninsurance rates, Medicaid enrollment rates, and cardiac hospitalization rates from the Behavioral Risk Factor Surveillance System; (5) HSA- and HRR-level measures of teaching hospital densities from the American Hospital Association;

and (6) state-level Part D plan enrollment and low-income-subsidy enrollment from CMS data. In total, we have 75 covariates that describe many of the differences across physicians, hospitals, and markets that are likely to be related to how physicians choose drugs, and how firms choose which physicians to interact with.

Table 3: Important Characteristics

	mean	s.d.	N selections
Own cardiov. claims	3,364	2,726	998
HSA Cardiol. annual claim avg. cardiov.	1,453	7,436	981
HSA Uninsured share	0.11	0.04	976
Hosp. admissions	17,591	17,538	968
HRR Share faculty	0.03	0.02	964
Zipcode-drug local market ad spend	33,716	48,374	953
HRR Num. AAMCs	16	15	922
HSA Medicare Advantage eligbl.	118,895	198,642	917
HSA Doc. annual claim avg. total	950	11,974	902
HSA Doc. annual claim avg. cardiov.	434	5,086	895
HRR Num. doctors	4,651	4,628	863
Hosp. Doc. annual claim avg. cardiov.	724	475	859
HSA Cardiol. annual claim avg. cardiov.	1,454	7,436	840
HRR Num. cardiol.	143	152	804
HSA Medicaid Share	0.22	0.08	795

Notes: Reports summary statistics for a subset of individual, hospital, and regional characteristics that are most frequently selected by the Lasso routine. N selections reports the number of Lasso estimations that select the variable out of 1,000 possible (2 Lassos for each of 500 bootstrap subsamples). For complete summary statistics, see Tables A5 and A6.

Table 3 reports the summary statistics for a select set of "important" covariates – those most frequently selected by the Lasso routine described in Section 4.2.3 below as determinants of payments or prescribing. Appendix G.1 reports the summary statistics for all covariates, along with the results from univariate regressions of our utilization and meal payment variables on each covariate. Notably, the associations between each covariate and these two outcomes often have different signs – many features that predict an increased likelihood of having a meal-based relationship also tend to predict a lower level of utilization.

This pattern, as well as the similarity of the prescribing patterns of paid and unpaid physicians shown in Figure 1 Panel (c) presents a puzzle. It suggests either that firms are spending large amounts of money on seemingly sophisticated efforts to influence physicians with no ultimate effect, or that those sophisticated efforts target and influence physicians who otherwise would have had low shares of the focal drugs. The rest of the paper is devoted to uncovering whether payments indeed have any effect on prescribing, the size (and distribution) of any such effect, and its welfare implications.

3 Regional Spillovers from Conflict of Interest Policies

This Section outlines our strategy to recover the causal effects of payments from firms to physicians. First, we outline our motivating intuition. Then we illustrate how our instruments affect the allocation of meal-based payments. Finally, we document patterns that support our exclusion restriction based on a series of placebo tests where we find particular subsamples for whom both the "first stage" (regressing payments on the instruments) and "reduced form" (regression utilization measures on the instruments) relationships are shut down.

3.1 Regional Policy Spillover Approach

To identify cardiologists who receive meals for plausibly exogenous reasons, we exploit the fact that, during the period we study, Academic Medical Centers (AMCs) across the U.S. had a wide range of policies that intended to prevent conflicts of interest (CoI) by limiting physician-firm relationships. We hypothesize that these CoI policies decreased the likelihood of physician-firm interactions not only for AMC faculty members, but for cardiologists who happened to have practices located nearby these institutions due to regional economies of scale in sales force allocation. This strategy is closely related to research designs recently employed in other industrial organization studies of sales (Hastings et al. 2017), and akin to a broader literature that examines behavior of bystanders exposed to externalities driven by aggregate features of their region (Waldfogel 2007).

The intuition of this approach is that drug firms, directly or via their marketing contractors, typically first determine marketing budgets and strategies based on aggregate market characteristics. Then the firms' "boots-on-the-ground" representatives use data analysis and their own knowledge of specific physicians to target high-value individuals.

Firms' marketing models can be very detailed and data-driven, and pharmaceutical sales forces maintain rich databases on prescribers' practice characteristics, prescribing behavior, and history of interactions with the firm (Campbell 2008). Thus it seems safe to assume that firms at least attempt to target physicians based on the expected incremental costs and benefits of sales effort. The expected benefit of interacting with a given physician depends on the size and appropriateness of the physician's patient panel, the physician's latent preferences over substitute products, and the physician's expected responsiveness to the payment and interaction.[15] Costs include the labor costs of additional sales representatives, the opportunity costs of diverting sales effort from other physicians, and any direct costs of the interaction

[15] In a notable recent example, Alpert et al. (2019) document that Purdue Pharma avoided marketing OxyContin in states with strict prescription drug monitoring programs.

(e.g., meal expenditure). They also implicitly include factors that limit or prohibit access for sales representatives. For example, the consulting firm ZS Associates publishes the *Access Monitor*^{TM} survey, which tracks and analyzes pharmaceutical representative access to physicians. The 2015 *Access Monitor*^{TM} report notes several key factors restricting access: academic medical centers' restrictive access policies, specialty-specific physician employment by hospitals and health systems that have central purchasing or otherwise limit physicians' autonomy, pressures on physicians that limit available time for firm interaction, etc. (Khedkar and Sturgis 2015).

Pharmaceutical sales territories are defined by geography and other organizing principles, such as therapeutic area (Campbell 2008). Given the fixed costs of deploying a sales force to a market, individual physicians' interactions with pharmaceutical firms will experience spillover effects from market-level characteristics. Thus, conditional on variables that proxy for individual physicians' attractiveness to pharmaceutical representatives—which may be correlated with physicians' underlying preferences—variables that proxy for attractiveness of *other* physicians in the same geographic market could be useful instruments for payments.

The AMC CoI policies that we focus on are described in detail in Larkin et al. (2017). We rely on data on AMCs' conflict of interest policies from the American Medical Student Association's (AMSA) conflict of interest scorecard. The AMSA scores evaluate the strictness of AMC policies regarding physician interactions with pharmaceutical/device companies, including salesperson access to AMC facilities, gifts to physicians, and enforcement of the policies.[16] We hypothesize that regions where AMCs have strong conflict-of-interest policies, as captured by high AMSA scores, will see fewer meal payments to physicians overall, even to those physicians unaffiliated with the AMCs. We further hypothesize that these effects will be stronger when a larger portion of the region's cardiologists are affiliated with the AMC and for cardiologists located more closely to the AMC in geographical space.

For the faculty in our sample (roughly 10 percent of cardiologists), the numerical conversion of these scores range from 0 to 31, where zeroes are assigned to AMCs that do not report any CoI policies to AMSA. To get a sense of how we use these scores, consider Sioux Falls, SD and Lubbock, TX, two cities whose major HRRs surround moderately sized state

[16]In every school year since 2007, medical schools have been asked to submit their policies to the AMSA for rating. Each institution's policy is graded in 13 different categories, including Gifts, Consulting, Speaking, Disclosure, Samples, Purchasing, Sales Reps, On-Campus, Off-Campus, Industry Support, Curriculum, Oversight, and sanctions for Non-Compliance. For each category except Oversight and Non-Compliance, the institution is assigned a numerical value ranging from zero to three. A zero is awarded if the institution did not respond to requests for policies or declined to participate; a one if no policy exists or the policy is unlikely to have an effect; two if the policy represents "good progress" towards a model policy; and a three if the policy is a "model policy." We generate aggregate AMSA scores for each institution; this aggregate ranges from 11 to 31-32 in 2011–2012.

universities with associated AMCs: the University of South Dakota and Texas Tech., respectively. These markets each have 24-28 cardiologists in our sample. However, the University of South Dakota's AMSA CoI score is 30 (95th percentile), vs. only 24 (25th percentile) at Texas Tech.

Also, many more of the cardiologists in the Sioux Falls region are faculty than in the Lubbock region. These differences are associated with large differences in meal rates: 16 percent in the Sioux Falls region vs. 41 percent in the Lubbock region. Likewise, we see large differences in prescribing: in Sioux-Falls, branded statins account for 2.1 percent of all cardiovascular prescriptions versus 2.9 percent in Lubbock. Of course, there may be other important differences in Sioux Falls vs. Lubbock that we want to account for, including the illness of the patient population, insurance rates, managed care penetration, and so on. This motivates our emphasis on physician-specific instruments and inclusion of rich control variables as we move to the full sample estimation strategy.

3.2 Effects of Policy Spillovers on Payments

Moving to the full sample, Figure 2 helps illustrate the relationships between meal receipt and different measures of AMSA CoI scores. The three subfigures at the top show raw relationships, conditioning only on molecule fixed effects. Each are scatterplots binned at percentiles of the data to illustrate the distribution of the policy scores. Panel (a) shows the relationship between AMC policies (as captured by the AMSA scores) and meal payments to cardiologist faculty affiliated with the AMC.[17] As expected, faculty at AMCs with more stringent policies are less likely to receive meal payments.

Panel (b) examines how these AMC policies spill over to affect non-faculty physicians working at the same hospitals as faculty. The horizontal axis is the AMSA score, weighted by the fraction of cardiologists at the hospital who are faculty, in order to capture the relative intensity of exposure that would logically enter a pharmaceutical firm's allocation of sales resources. The figure shows that non-faculty physicians are less likely to receive meals if they work at hospitals with more faculty affiliated with AMCs with restrictive policies.

Panel (c) shows a similar pattern of spillovers at the HRR level, and Appendix G.2 shows a similar pattern at the HSA level. Cardiologists are less likely to receive meal payments from AstraZeneca and Pfizer if they work in regions where more cardiologists are affiliated with AMCs with more restrictive policies, even though those policies do not directly govern the focal cardiologists' own or own affiliated hospitals' behavior.

These relationships are consistent with our conversations with current and former phar-

[17] The faculty linkage is from the Association of American Medical Colleges (AAMC) faculty roster.

Figure 2: AMSA-scored Conflict-of-Interest Policies and Meals

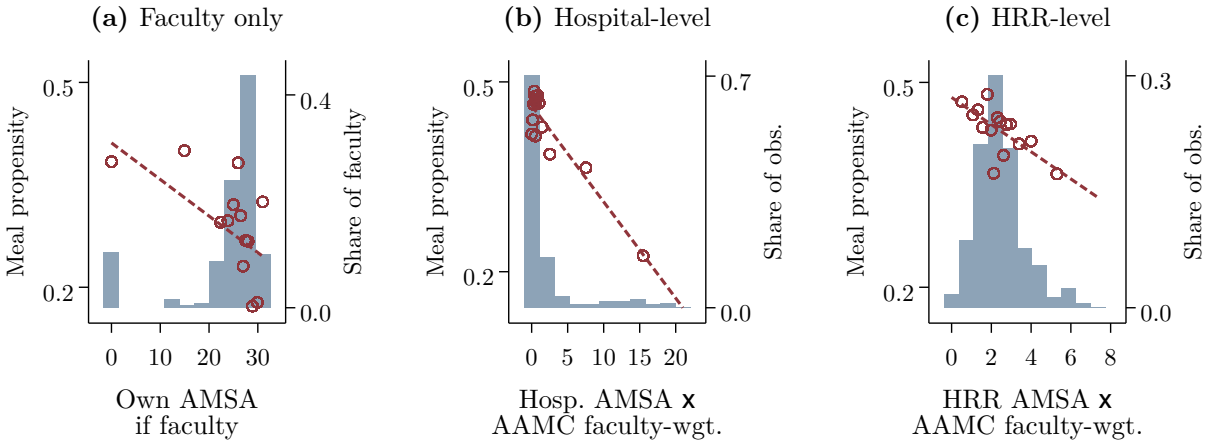

(d) I.V. and Interaction-Term Summary Statistics

	$p10$	$p50$	$p90$	mean	s.d.
HRR AMSA	24.1	25.8	27.6	25.6	3.2
Time to nearest AAMC (mins)	1.0	5.8	40.2	14.6	31.0
HRR AAMC faculty wgt. (share)	0.010	0.027	0.050	0.029	0.018

(e) HRR-level I.V. First Stage

	(1)	(2)	(3)	(4)
HRR AMSA	–0.0368	–0.0357	–0.0314	–0.0299
	(0.0033)	(0.0032)	(0.0028)	(0.003)
HRR AMSA ×		0.0083		0.0089
time to nearest AAMC		(0.0041)		(0.0042)
HRR AMSA ×			–0.0265	–0.0267
AAMC faculty wgt.			(0.0016)	(0.0016)

Notes: (a-c) display equally binned scatterplots of the unconditional correlation between meals and three different AMSA score metrics from (a) the cardiologists' own AAMC (if faculty), (b) the faculty-weighted AMSA score of a cardiologist's hospital, or (c) the faculty-weighted AMSA score of a cardiologist's HRR, excluding the scores of faculty within their own HSA and hospital. The lefthand axis is for the scatterplots and linear fit lines, while the underlying histograms of the different scores are described by the righthand axis. (d) reports summary statistics for key variables used as instruments or interactions with instruments. (e) reports the OLS regressions of meal indicator on the vector of Lasso-selected controls (own- and hospital-level AMSA scores are potential controls) and manually selected HRR-level instruments. All instruments are standardized. Point estimates and standard errors use the bootstrap approach as described in Section 4.2 and Appendix D. The average number of observations per bootstrap sample is 26,045.

maceutical sales executives and pharmaceutical marketing consultants regarding economies of scale in sales force allocation. The primary concern with using these AMC policies as instrumental variables is that the exclusion restriction may fail due to direct effects of conflict of interest policies on norms regarding prescribing, or due to unobservable factors correlated with selection into more restrictive policies (see discussion in Larkin et al. 2017). We take several steps to refine and test our IV strategy in order to mitigate this concern.

One refinement we make is to focus on spillovers to those physicians who are more likely to be "bystanders" to an AMC's policies. We use the own AMSA score (for faculty) and own hospital AMSA score (for those working at hospitals with faculty) as control variables, and we only use the jackknifed HSA and HRR spillover measures as instruments. This focuses on spillovers to cardiologists who have less direct ties to the AMC or its affiliated physicians, mitigating the concerns of the spillover being correlated with unobserved determinants of prescribing.

The second refinement we make is to use: (1) the (spillover) AMSA score itself, (2) variables interacting the AMSA score with the percent of cardiologists affiliated with the AMC (as in the plots above), and (3) interactions with the drive time of the cardiologist's office from the AMC (see Panel (d) of Figure 2 for summary statistics). The first stage relationship between these interactions and whether the focal physician receives a meal provides further evidence related to the economies of scale hypothesis. Panel (e) of Figure 2 shows results for regressions with molecule fixed effects and a large vector of control variables.[18] Having an AMC with a higher AMSA score in the region is negatively associated with receiving a meal. The coefficient magnitudes indicate that a 1 standard deviation increase in the HRR-based AMSA score causes a roughly 3 p.p. decline in the probability of receiving meals from a firm (relative to a base meal probability of approximately 0.5). This effect is stronger if a larger percentage of cardiologists in the region are associated with the AMC (see the second row of Panel d), and the effect is weaker if the focal physician's office is located further from the AMC (see the third row of Panel d).

3.3 Placebo Tests of the Exclusion Restriction

While we cannot formally test the exclusion restriction directly, we explore its validity by conducting a set of placebo tests in the spirit of Angrist et al. (2010). The motivating logic is that, in a reduced form regression of the outcome variable on an instrument, the coefficient on the instrument will depend on the causal relationship, plus any potential correlation with unobservable determinants of prescribing (the presence of which would indicate a violation

[18] Below in Section 4 we describe this large set of covariates and the Lasso-based selection procedure we use to choose which covariates to include as controls in these regressions.

of the exclusion restriction). For a subsample of physicians whose observable characteristics make them very unlikely to receive a meal payment, the causal relationship is substantially shut down. Thus, for such a sample, a significant coefficient on the instrument in the reduced form would be evidence of a violation of the exclusion restriction. Conversely, a zero coefficient provides evidence consistent with the exclusion restriction being satisfied in the subsample, which in turn provides some reassurance that it holds in the full sample.

We conduct several tests based on this logic, the results of which are in Appendix G.3. Briefly, in four distinct subsamples—cardiologists in states with restrictions on meals, faculty cardiologists facing strict CoI policies, cardiologists at hospitals with strict CoI policies, and cardiologists whose observables make them unlikely meal targets—the first stage relationship shrinks by an order of magnitude, and the reduced form relationship disappears. This is encouraging evidence consistent with the exclusion restriction.

4 Effects of Meals on Demand for Statins

This Section develops an explicit demand model of how consumers trade off the influences of meals and out-of-pocket prices and substitute across competing drugs. This allows us to estimate an effect of meal payments, conditioning on other drivers of demand. Our model also allows for heterogeneity in the treatment effects of payments across physicians. These estimates are of direct interest in better understanding how firms target and influence physician intermediaries. They are also an important input into our welfare analysis in Section 5, where we add a supply-side drug pricing model and allow for the possibility that physician/patient decision making could be imperfect due to physician agency or physician/patient informational or behavioral biases.

4.1 Model of Demand for Pharmaceuticals

Let the utility of molecule $j \in \mathcal{J} = \{0, 1, ..., J\}$, for use case i (a doctor/patient/visit combination) in each market defined by doctor d in year t be given by: $u_{idjt} = \delta_{djt} + \varepsilon_{idjt}$.[19] The choice $j = 0$ represents the choice of treatment other than a statin, with mean utility normalized to $\delta_{d0t} = 0$. We measure the market size of potential statin patients as the number of all cardiovascular prescriptions, including other lipid-modifying drugs, for each physician-year. The use-specific i.i.d. unobservable $\varepsilon_{idjt} = \epsilon_{idt} + (1 - \lambda)\epsilon_{idjt}$ is the random coefficients representation of the nested logit model (Cardell 1997), where ϵ_{idt} is a random

[19] The only molecule sold in both branded and generic format during the time period we study is Lipitor / atorvastatin in 2012. They have different j indices, allowing preferences for the two to be potentially different and flexibly estimated.

component common to statins vs. alternative treatments, and ϵ_{idjt} is the standard type I extreme value error term (with scale normalized to one) that is i.i.d. across molecules. As the nesting parameter $\lambda \in [0, 1]$ approaches 1, there is less substitution outside the nest.

We specify mean utility across use cases as:

$$\delta_{djt} = \theta^m_{dj} 1_{\{m_{dj}>0\}} - \theta^p p^{oop}_{djt} + X_{djt}\theta^x_j + \xi_{djt} \ . \tag{1}$$

Here, $\theta^m_{dj} 1_{\{m_{dj}>0\}}$ is an indicator for whether provider d received a meal from the manufacturer and its utility weight. Importantly, this utility weight may be specific to the drug-doctor pair. It may even be negative and lead to decreased prescribing. This heterogeneity across drugs and doctors in the impact of a meal on mean utility captures several sources of variation that have been discussed in prior research (e.g., Inderst and Ottaviani 2012) such as: the informative vs. persuasive nature of the interactions associated with meals, physician prior knowledge/ability, physician concern for patients, and the fraction of patients that are wary/sophisticated/informed.

Turning to the other components of mean utility, $\theta^p p^{oop}_{djt}$ is the average out-of-pocket price paid by patients and its weight. $X_{djt}\theta^x_j$ is a rich set of covariates that captures perceived quality variation across molecules, as well as regional and provider variation in prescribing patterns over time (we return to discuss this in detail when we discuss estimation of the model in Section 4.2). Finally, ξ_{djt} is a product-physician-year unobservable preference heterogeneity term.[20]

Given a set of products available to a provider $\mathcal{J}_{dt}$ and flow of choice opportunities Q_{dt}, we assume the provider/patient chooses the product that maximizes decision utility, so that expected quantities demanded are given by:

$$q_{djt} = Q_{dt} Pr[u_{idjt} > u_{ikdt}, \forall k \in \mathcal{J}_{dt}] = Q_{dt} \frac{e^{\frac{\delta_{djt}}{1-\lambda}}}{\sum_{k\in\mathcal{J}_{dt}} e^{\frac{\delta_{kdt}}{1-\lambda}}} \frac{\left(\sum_{k\in\mathcal{J}_{dt}} e^{\frac{\delta_{kdt}}{1-\lambda}}\right)^{1-\lambda}}{1 + \left(\sum_{k\in\mathcal{J}_{dt}} e^{\frac{\delta_{kdt}}{1-\lambda}}\right)^{1-\lambda}} \ . \tag{2}$$

4.2 Demand Identification and Estimation

In this Section, we show how the role of meal payments can be fit into a standard potential outcomes framework, embedded in the structural demand system. The potential for selection both into treatment with a meal and on the (heterogeneous) response to treatment

[20] Several recent papers (e.g., Dubois et al. 2018; Shapiro 2018b; Sinkinson and Starc 2019) focused on television advertising have explicitly focused on the possibility that such ads can have spillover effects across brands in a category. In Appendix G.5 we explore and do not find strong evidence for such spillovers in our context, perhaps related to the more focused nature of meal payments to physicians vs. television advertisements to consumers.

suggest a semi-parametric instrumental variables strategy, using the data and institutional details discussed previously. We also need to identify and estimate the other parameters of the demand system, in particular the price and nest parameters, which face the common simultaneity/endogeneity problems inherent to any demand estimation exercise.

Our estimation approach proceeds in three broad steps, which we take for each of 500 bootstrap samples. We outline the strategy here and provide more specifics in the remainder of the Section. Appendix D provides further details. In the first step, we resample at the use-case level (approximating patient-level decisions) to account for sampling error in market shares and estimate the price and nest parameters and a set of molecule-doctor fixed effects. To account for the endogeneity of prices and nesting patterns, we use instrumental variables that leverage the panel variation in prices and choice sets induced by the introduction of generic atorvastatin. In the second step, we set up a potential outcomes framework where the molecule-doctor fixed effects are the outcome of interest and the key endogenous variable is the indicator for meal payments. We estimate this selection equation, which produces a propensity score for meal allocation across doctors. The third step then uses these propensity scores and our policy-spillover instruments to recover the distribution of marginal treatment effects across the sample of molecule-doctor observations. For each subsample, we also employ a jackknife approach (where the square root of the number of unique cardiologists (120) are dropped in each resampling), blocked at the cardiologist level to allow for arbitrary correlations within cardiologist.

As a preliminary step, we first linearize the demand model, following Berry (1994). We set choice probabilities implied by the demand model in Equation (2) equal to observed market shares, and invert the system of equations to obtain mean utilities as a function of the market shares: $\delta_{djt} = \ln(s_{djt}/s_{d0t}) - \lambda \ln(s_{dj|_g t})$. Combining this with Equation (1) yields the linear specification:

$$\ln(s_{djt}/s_{d0t}) = \lambda \ln(s_{dj|_g t}) + \theta^m_{dj} 1_{\{m_{dj}>0\}} - \theta^p p^{oop}_{djt} + X_{djt}\theta^x_j + \xi_{djt} \ . \tag{3}$$

where s_{djt} is j's overall market share, s_{d0t} is the market share of the outside good (non-statin treatments), and $s_{dj|_g t}$ is j's market share within nest g, the set of statin treatments.

4.2.1 Estimating Price Sensitivity and Nest Parameters

In the first stage of estimation, we implement a differences-in-differences style estimator to leverage the price and choice set variation resulting from the introduction of generic atorvastatin at the end of 2011. We find this to be the most compelling specification to

identify the coefficients on price and within-nest share. We estimate:

$$\ln(s_{djt}/s_{d0t}) = \lambda \ln(s_{dj|gt}) - \theta^p p_{djt} + \psi_{dj} + \theta_t + \theta_{\text{Lip12}} + \tilde{\xi}_{djt} \tag{4}$$

where ψ_{dj} is a product-doctor-specific fixed effect and θ_t is a year fixed effect. We further include θ_{Lip12}, a coefficient for Lipitor in 2012, to capture the fact that demand for branded Lipitor is small and idiosyncratic in 2012 as it is removed from formularies over the course of the year. With a slight abuse of notation, we use a single fixed effect for both branded Lipitor and generic atorvastatin in order to leverage the within-molecule variation in price between 2011 and 2012 induced by generic entry.

There are two immediate empirical difficulties associated with this specification. First, the ψ_{dj} terms are measured with noise. In order to account for this issue, we employ a version of the standard shrinkage approaches from the empirical Bayes literature (see Chandra et al. (2016) for a recent application in the healthcare context). We modify the standard approach by resampling at the "use case" level to account for sampling error in market shares. Appendix D.4 provides a detailed description of the procedure and illustrates how it adjusts the ψ_{dj} distribution.

Second, we must account for the likely endogeneity of $\ln(s_{dj|gt})$ and p^{oop}_{djt}. We leverage both the average changes induced by generic atorvastatin entry and also the heterogeneity in insurer responses to this entry across geography (described in detail in Appendix A.2). When Lipitor's patent expired, some insurers instantly added generic atorvastatin to their preferred drug lists and/or removed Lipitor from their formularies, while others took more than a year. The variation in penetration of these insurers across geography generated large variation in the relative prices consumers faced for Lipitor and generic atorvastatin. To utilize this variation, we create instruments for each plan-drug-year-region as the average out-of-pocket price for that drug-year-insurer across *other* regions. We then average across plans to create an instrument for physician d's region. We also create an analogous instrument based on an average dummy for formulary coverage. The instrument set is then: $Z^p = [p^{oop,IV}_{djt}, \bar{1}_{\{j \in form^{IV}_{dt}\}}]$. These are similar to the bargaining ability instruments in Grennan (2013, 2014) and Dickstein (2016), with the added benefit of a clear mechanism driving their variation. As such, they are also valid for both $\ln(s_{dj|gt})$ and p_{djt}.

In addition to the instruments linked to generic atorvastatin entry, we also follow much of the literature (e.g., Berry and Waldfogel 1999) in using a polynomial in the cardinality of the sets of statins and strong statins prescribed $Z^J = [\ln(|\mathcal{J}_{dt}|), |\mathcal{J}_{dt}|, |\mathcal{J}_{dt}|^2, \ln(|\mathcal{J}^{ss}_{dt}|), |\mathcal{J}^{ss}_{dt}|, |\mathcal{J}^{ss}_{dt}|^2]$ as instruments. This leverages the fact that more variety will mechanically affect within-group shares. In this particular context, it is also related to the intuition behind Sinkinson

and Starc (2019), who use managed care penetration to proxy for restricted choice sets in the statin market in an earlier time period.

4.2.2 Estimating the Effects of Meals on Prescribing

The fixed effects ψ_{dj} from the first step of our estimation capture all of the sources of persistent prescribing differences across doctors during our sample period. We now turn to recovering the determinants of this prescribing heterogeneity, in particular, the extent to which it is influenced by meal payments from pharmaceutical firms. To do this, we project the product-doctor fixed effects on our cross-sectional meal indicator and a rich set of controls for physician and market characteristics.[21]

$$\hat{\psi}_{dj} = \theta^m_{dj} 1\{m_{dj} > 0\} + \theta_j + X_{dj}\theta^x_j + \bar{\xi}_{dj} \ . \tag{5}$$

In our preferred specification, we construct $1\{m_{dj} > 0\}$ as a dummy for physician d receiving any payment from Pfizer over 2010-2012 (in the case of j=Lipitor/atorvastatin), or as a dummy for physician d receiving any payment from AstraZeneca over 2011–2012 (in the case of j=Crestor).[22] We estimate this equation only on observations for Lipitor/atorvastatin and Crestor, as generic firms do not provide meals to doctors.

The focal parameter θ^m_{dj} describes the effect of industry interaction on a physician's use of branded statins. We suppose that there is an underlying model of firms allocating meals to doctors as a function of the doctor-specific return on investment and regional economies of scale. Appendix C.2 presents an explicit characterization of such a model. Meal decisions are likely based on data we have available as researchers, plus other factors that are unobservable to us. This induces the potential for selection on unobservables, both in levels and in terms of the response to the meal "treatment." In order to accommodate this in our estimation, we also specify a selection equation that is a semi-parametric representation of a model of strategic meal allocation.[23] This selection equation takes the form of a linear probability model:

$$1\{m_{dj} > 0\} = Z_{dj}\gamma^z_j + X_{dj}\gamma^x_j + \mu_{dj} \ . \tag{6}$$

The outcome equation (5) and selection equation (6) fit into the canonical potential outcomes framework. In the context of the model, the unobservable in the selection equation μ_{dj} may be correlated with both $\bar{\xi}_{dj}$ and the heterogeneous component of θ^m_{dj}. In such a

[21] The idea of a secondary regression to uncover the determinants of fixed effects goes back at least to Mundlak (1978).

[22] Payments from AstraZeneca in 2010 are not available in our data.

[23] Appendix C.2.1 shows the tight relationship between the selection equation here and a structural version of the model in Equation (10) for a particular cost function with increasing returns to scale.

case, the standard 2SLS estimator will be a particular weighted average of the local average treatment effects on compliers, will be higher than the average treatment effect, and will have limited relevance for policy simulations (Andresen 2018). We thus estimate the marginal treatment effects directly. We can then estimate many treatment effects of interest as a function of the distribution of MTEs.

We take a semi-parametric approach to MTE estimation. We first estimate the propensity to receive a meal using the linear probability model in (6). We then use the residuals from this estimated propensity score model (which the literature typically refers to as the "unobserved resistance to treatment") and a nonparametric instrumental variables approach to map out the MTEs across observations with similar unobserved resistance to treatment. For a detailed discussion, see Appendix D.2.

4.2.3 Lasso Selection of Controls and Instruments

The cross-sectional nature of our identification strategy and the data-intensive nature of our semi-parametric MTE estimation strategy make a rich set of controls and flexible functional forms especially important. Relatedly, we have no a priori theory for the functional form relating our potential instruments to meals. To address these issues, we include the large set of potential controls at the regional, hospital, and doctor level discussed in 2.5. We include linear, quadratic, and logarithm terms of each variable. We also include the potential instruments based on AMSA scores of local AMCs and their interactions with the relative presence of the AMC in the region and geographic location of each doctor. Finally, we include linear, quadratic, and logarithm transformations of the instruments, and further interact them with dummy variables for each manufacturer.

As we allow the control and instrument sets to grow larger and more flexible, we run into the issues of sparsity and collinearity which have been the topic of a growing literature at the intersection of econometrics and machine learning. We follow Belloni et al. (2017); Chernozhukov et al. (2018) and related literature in using Lasso to select the controls and instruments which most strongly predict meals and prescribing. Appendix D.1 discusses the details of this procedure. Briefly, we use the split-sample cross-fitting approach of Chernozhukov et al. (2018) – tuning the Lasso model (using 10-fold cross validation) on one half of the data and estimating the model parameters on the other half – to efficiently remove bias due to overfitting. The final column of Table 3 reports, for the subset of variables that were most frequently selected by the Lasso routine, the number of Lasso estimations that selected each variable (out of 1,000 possible).

4.3 Results: Price Coefficients and Substitution Patterns

Table 4 provides estimates of the first step of the demand model, using the panel variation from generic atorvastatin entry to recover price and nest parameters as well as molecule-doctor fixed effects. Column (3) details the estimates from our preferred model, which match basic institutional details and prior literature. The other columns illustrate the importance of using molecule-doctor fixed effects (these are absent from column (1) and present in column (2)) and our instruments (employed in column (3) only)—without them, the demand estimates imply an extremely small or even positive relationship between demand and price.

Table 4: Demand Estimates Step 1—Price and Nest Coefficients and $\{\psi\}$

	(1) OLS, ψ_d	(2) OLS, ψ_{dj}	(3) I.V., ψ_{dj}
θ^p	0.00110	–0.00020	–0.00761
	(0.00002)	(0.00002)	(0.00017)
λ	0.944	0.972	0.421
	(0.0004)	(0.001)	(0.011)
mean(η^p)	0.292	–0.103	–0.211
s.d.(η^p)	0.329	0.116	0.228
F–stat.			680.1
$R^2(\delta_{djt} : \psi_j)$			0.282
$R^2(\delta_{djt} : \psi_j - \theta^p p)$			0.417
$R^2(\delta_{djt} : \psi_{dj} - \theta^p p)$			0.802
mean($\psi_{dj}/\lvert\theta^p\rvert$) strong statins			–294.4
mean($\psi_{dj}/\lvert\theta^p\rvert$) other generics			–313.3
s.d.($\psi_{dj}/\lvert\theta^p\rvert$)			74.5

Notes: Reports parameter estimates from Eq. 4. Point estimates are based on the average of the 500 perturbed-bootstrap samples; standard errors for the main parameters (θ^p and λ), in parentheses, are based on the standard deviation of these 500 point estimates. $R^2(\delta_{djt} : \cdot)$ reports the R^2 from a regression of $\delta_{djt} \equiv \ln(s_{djt}/s_{d0t})$ on some combination of the molecule- (ψ_j) or molecule-doctor-level (ψ_{dj}) fixed effects, and possibly the price effect ($\theta^p p$). The average number of observations in each bootstrap sample is 116,559.

The nesting parameter estimate of 0.421 in our preferred model indicates that there is more substitution within statins than to non-statin alternatives. This matches institutional knowledge that there are certain types of cardiovascular patients for whom statins are appropriate. The price coefficient is small but nontrivial, as we would expect given the muted incentives implied by insurance, and the related own-price elasticity $\eta^p = \frac{\partial s}{\partial p}\frac{p}{s}$ of -0.21 is within the range of prior estimates for the Part D setting (e.g., Abaluck et al. 2018; Einav et al. 2018). In Appendix G.4, Table A9 reports a range of alternative specifications that illustrate the importance of instrumenting both the price and nest parameters, and that also show similar results if we drop AMC faculty, do not perform the quantity-perturbation

procedure, or use alternative nesting structures.

The middle panel of rows in Table 4 provide useful summary statistics on the distribution of physician-molecule fixed effects recovered (ψ) and the role of these and other factors in explaining variation in prescribing behavior in the data.[24] The first three rows give a sense of the variation in mean utilities δ_{djt} explained by molecule fixed effects ψ_j (28 percent), price (an additional 13 percent), and physician-molecule fixed effects ψ_{dj} (an additional 39 percent). Thus, price does indeed impact demand, but it explains less of the variation than molecule or physician-molecule specific factors. The large amount of additional variation explained by ψ_{dj} relative to ψ_j indicates that physician-specific preferences within molecule explain more variation in prescribing than the average differences across molecules themselves. The differences across molecules are summarized in the next two rows in the mean fixed effects (scaled into dollars) for the strong statins versus the other generic statins. On average, cardiologists value the strong statins about \$19 more than the generics, which is in line with the observed OOP prices (in 2011, the branded strong statins' OOP was about \$27 more than the generics' OOP). The overall physician-molecule preference variation itself is large, with one standard deviation equivalent to an OOP price differential of about \$75.

4.4 Results: Meal Payments and Prescribing

The second part of our estimation procedure considers the relationship between meal payments and the estimated physician-molecule preference parameters ψ_{dj}. This section summarizes the results of that procedure and several key economic implications of those results.

Figure 3, Panel (a) shows histograms of the first stage propensity score estimates for the physician-molecule observations with and without meal payments. The model produces large overlapping support for the two groups across the unit interval. In our application, we have found that the rich specification of controls and instruments enabled by the Lasso approach is critical to achieving both this rich overlapping support, and also sufficient variation in the instruments conditional on the propensity score. Appendix D.3 provides more details on the most frequently selected controls and instruments from our Lasso routines as well as the importance of a rich specification for the first stage.

Figure 3, Panel (b) plots our MTE estimates vs. the unobserved resistance to treatment.[25] The average treatment effect of 0.30 is roughly equivalent to a \$40 price decline, but the distribution of estimates rejects the hypothesis of a homogenous treatment effect. At the

[24] See Appendix D.4 for more on how the quantity perturbation and empirical Bayes shrinkage procedures influence the estimates of ψ_{dj}.

[25] The literature on MTE estimation defines the unobserved resistance to treatment as the quantiles of the distribution of residuals from the first stage propensity score estimation.

Figure 3: Meal Payment Propensity and Marginal Treatment Effects

(a) Pr(Meal | X,Z)

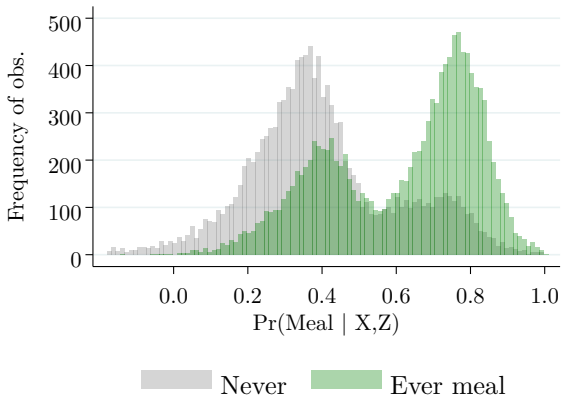

(b) MTE Plot

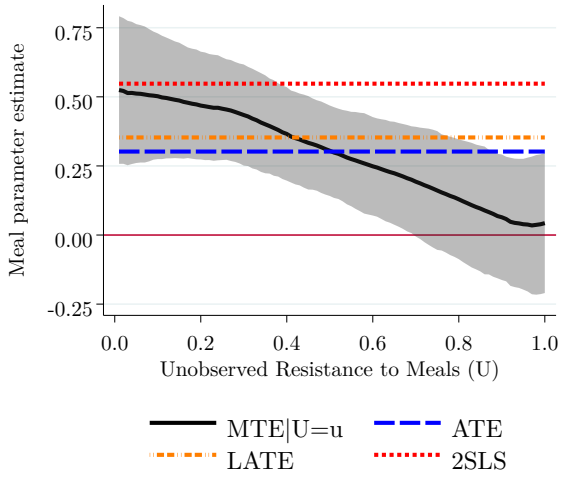

(c) MTE Point Estimates

			MTE-based	
	OLS	2SLS	LATE	ATE
	(1)	(2)	(3)	(4)
θ^m	0.089	0.528	0.344	0.300
	(0.0013)	(0.1188)	(0.0824)	(0.0855)
F-stat			138.8	

Notes: (a) plots the distribution of LPM-based propensity scores, by actual meal status. (b) plots the MTE curve ($\mathbf{E}[\theta^m \,|\, U = u]$) with 95 percent C.I. in shaded grey, alongside other point estimates of θ^m. (c) displays the estimates from the OLS (Col. 1), 2SLS (Col. 2) and MTE (Cols. 3–4) specifications, all of which include the Lasso-selected controls. The average number of observations in each bootstrap subsample is 26,045.

10th percentile of unobserved resistance (i.e., physicians that firms are very likely to pay), the effect is 0.51 (equivalent to a $67 price decline or a 1.0 standard deviation increase of the underlying physician-molecule preference heterogeneity), while the effect at the 90th percentile of unobserved resistance (i.e., physicians that firms appear to avoid), the point estimate is 0.06 (a $9 price decrease or 0.13 standard deviation increase in the underlying heterogeneity) and it is not statistically distinguishable from zero. Appendix G.6 shows that the shape of these MTE estimates is largely robust to alternative samples and modeling decisions.

The table at the bottom of Panel (b) compares several estimates of θ^m: ordinary least squares (column (1)), two-stage least squares (column (2)), and the LATE (column (3)) and ATE (column (4)) associated with the marginal treatment effects (see Andresen (2018) Table 3 for comparison of the weights in 2SLS and MTE-LATE). As suggested by Heckman et al. (2006), the 2SLS estimator overestimates the ATE in this case with positive selection on heterogeneous treatment effects. Our estimated average treatment effects are larger than those found in other papers that address physician selection into receiving payments with the inclusion of physician fixed effects (Agha and Zeltzer 2019; Carey et al. 2020; Shapiro 2018a). For example, Shapiro (2018a) finds that a detailing visit increases prescribing of antipsychotics by 14 percent in the subsequent year,[26], whereas the coefficients in our nested logit demand model imply that a meal-based detailing relationship increases promoted statin prescribing by 45 percent for the average physician, but by 63 percent for the average physician actually targeted by firms and only 27 percent for physicians firms avoid.[27] Interestingly, the OLS estimate is smaller than the ATE, suggesting that firms target physicians that would have otherwise prescribed relatively low shares of strong statins, thus muting the correlation between meals and strong statin shares.[28]

An advantage of the MTE estimation approach is that the resulting estimates can be paired with the data (i.e., physician features and realized treatment) to derive the expected response to treatment $E[\theta^m_{dj}]$ for any observation in the data.[29] We find these and other economic parameters more intuitive to interpret, and so we move to them to discuss the

[26]We credit Carey et al. (2020) for this calculation.

[27]The effect of the overall relationship may be much larger than the within-physician effect of an incremental meal. Chintagunta and Manchanda (2004); Shapiro (2018a); Agha and Zeltzer (2019) each consider the role of detailing "stock." Agha and Zeltzer (2019) also explicitly focus on diffusion of drugs at the beginning of their life cycles.

[28]These results stand in contrast to papers that model total quantity or expenditure as the dependent variable, which generally find that addressing physician selection (often using doctor fixed effects) results in smaller estimates of meal effects. Prior research has noted that firms target physicians with large relevant patient panels. See, e.g., Chintagunta and Manchanda (2004).

[29]More formally, $E[\theta^m_{dj} \mid X_{dj}, 1\{m_{dj} > 0\}]$; see Appendix D.2, and Eq. 16 specifically, for more on how individual-level expected treatment effects are derived from the MTE model.

treatment heterogeneity.

Figure 4, Panel (a) summarizes the distribution of expected treatment effects, normalized by the standard deviation of the physician-molecule preference variation, $E[\theta^m_{dj}]/SD(\psi^m_{dj})$. Payments are clearly directed to physicians with higher expected response to treatment. The histogram for those who received meal payments is shifted far to the right of the one for those who did not receive treatment. The average expected response of those receiving payments is a 0.72 standard deviation change in mean prescribing preference for the focal molecule. By contrast, for those not receiving payments, the same average effect is 0.30, and for about two-thirds of these not-paid cardiologists the effect is not statistically different from zero. The difference between the centers of these distributions is driven to a great extent by the steepness of the gradient documented above in Figure 3 panel (b), which implied a sizable difference between the average treatment effect on the treated and the average treatment effect on the untreated.

Panel (b) of Figure 4 plots a similar set of histograms for a different variable of interest, the expected prescribing share of the focal drug with no meals, $E[s_{dj}(m_{dj} = 0)]$. This is a partial equilibrium counterfactual computation, and it helps to solve the puzzle of why the summary statistics showed no clear difference between prescribing patterns for physician-molecule observations with and without meals. Here, the histogram for those receiving meal payments is shifted to the left of those who do not, indicating that meal payments tend to go to physicians who would have otherwise prescribed below average amounts of the focal molecule. On average, the effect of meals is to bring prescribing patterns by those who receive meals into line with those who do not. While this is indirect evidence, it does seem consistent with a story of meal payments (and the interactions surrounding them) providing information or reminders that counteract potential underprescribing for some physicians.

Panel (c) of Figure 4 plots the distribution of expected changes in revenue from targeting meal payments for each physician-molecule, bringing together several of the important dimensions of meal targeting—selection on patient volume, selection on expected response, and selection on expected counterfactual prescribing patterns—into one measure. The distribution for treated physicians is shifted significantly rightward from that of untreated physicians. Meals increased prescribing revenue to drug firms by roughly $4,900 for the average treated physician. However, counterfactually extending meals to all untreated cardiologists would have increased prescribing revenue to drug firms by only about half of that.[30] When compared to an estimated average cost of detailing physicians of about $1,600-$1,800 per physician-year (Liu et al. 2020), these estimates provide a large piece of the puzzle regarding

[30]In quantity terms, meals increase promoted statin prescribing by 45 percent for the average physician, by 63 percent for the average paid physician, and by 27 percent for the average physician not paid.

Figure 4: Heterogeneity in Treatment Effects by Observed Meal Receipt

(a) $E[\theta^m_{dj}]/SD(\psi_{dj})$

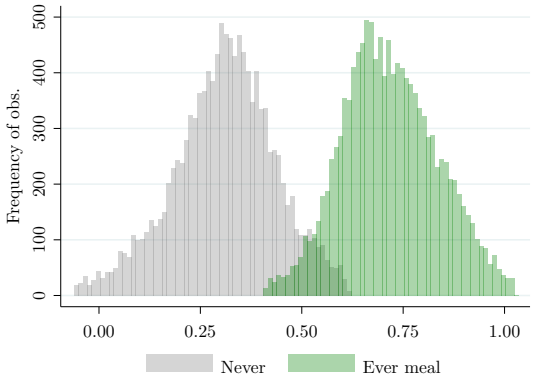

	$p10$	$p25$	$p50$	$p75$	$p90$
Never	0.108	0.214	0.316	0.416	0.507
	(0.175)	(0.175)	(0.178)	(0.175)	(0.181)
Ever	0.541	0.621	0.713	0.823	0.921
Meal	(0.174)	(0.181)	(0.180)	(0.182)	(0.182)

(b) $E[s^{m=0}_{dj}]$

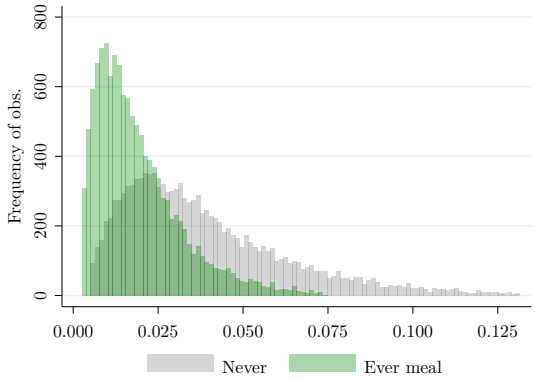

	$p10$	$p25$	$p50$	$p75$	$p90$
Never	0.013	0.021	0.033	0.052	0.075
	(.)	(.)	(.)	(.)	(.)
Ever	0.006	0.010	0.016	0.026	0.038
Meal	(0.001)	(0.002)	(0.003)	(0.004)	(0.006)

(c) $p_{dj} \times E[q^{m=1}_{dj} - q^{m=0}_{dj}]$

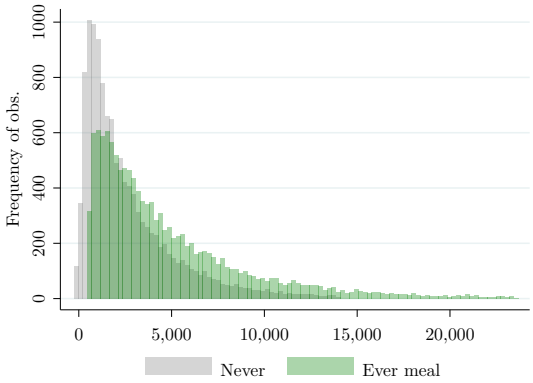

	$p10$	$p25$	$p50$	$p75$	$p90$
Never	346.8	814.5	1,809.1	3,705.2	6,635.1
	(556.1)	(725.4)	(1,221.0)	(2144.8)	(3,411.4)
Ever	1,019.1	1,796.7	3,500.3	6,549.9	11,083.2
Meal	(206.4)	(349.1)	(654.6)	(1,201.0)	(1,892.8)

Notes: Plots the posterior estimates of the dj-specific estimates of meal effects or counterfactuals. (a) Plots the distribution of meal parameters (θ^m) scaled by the standard distribution of the doctor-molecule mean utility "fixed effects" (ψ_{dj}). (b) Plots the (partial equilibrium) counterfactual distribution of cardiovascular shares if no doctors received meals. (c) Plots the distribution of marginal revenues due to meals. Revenues are based on our estimate of the price p^{mfr} paid to manufacturers, net of rebates and markups charged by supply intermediaries. As discussed in Appendix E, we assume $p^{mfr} = 0.55 * p^{pos}$ in the main text, based on recent estimates of average rebates and markups for cardiovascular and branded drugs. However, our takeaways are qualitatively similar if we assume a higher pass-through rate of $p^{mfr} = 0.68 * p^{pos}$. Beneath each plot are the point estimates and standard errors for select percentiles of these distributions by treatment status.

why some physicians are targeted by firms and others are not.[31]

5 Equilibrium Welfare Effects of Meals

The above results demonstrate that meals have large effects on prescribing, and that they are targeted to particularly responsive physicians who would otherwise prescribe below-average amounts of branded firms' drugs. However, the equilibrium effect of meal payments from pharmaceutical firms to physicians also depends upon how they interact with distortions from other market imperfections. In this Section, we use the results from the structural drug demand analysis, together with a model of drug pricing, to investigate the impact of a counterfactual meal ban on equilibrium prices, quantities, and welfare. Our framework includes a parameter governing the extent to which meals improve prescribing in the presence of possible underlying decision errors, allowing us to map demand into welfare for a range of assumptions regarding the optimality of baseline prescribing behavior. Finally, we use estimates regarding health outcomes from the clinical literature to calibrate decision errors in the context of our model and discuss the implications of that exercise for welfare.

5.1 Graphical Framework

To build intuition regarding these points, consider the welfare effects of payments that shift the demand curve outward. Panel (a) of Figure 5 presents a hypothetical demand curve in blue and a "biased" demand curve shifted outward in red. Assuming without loss of generality that the drug's marginal cost is zero, the welfare loss under perfect competition is shown in the shaded triangle below the line segment $\overline{Q^{eff}Q^{b}}$—marginal patients prescribed the drug in the presence of payments to physicians receive negative health benefits.

In a setting with perfect competition, this conceptual framework suggests that a simple analysis of the causal effects of payments on prescribing is all that is needed. However, in many empirically relevant settings with firm payments to experts, firms also have market power, and utilization is distorted away from the social optimum due to high prices. In prescription drug markets, branded products have patent protection, and they often compete with differentiated branded and generic substitutes whose manufacturers make their own strategic pricing and promotion decisions. Payments are typically only made for branded drugs as generic margins are too small to justify such costly marketing. A simple version

[31] Liu et al. (2020) estimate that Pfizer (AstraZeneca) visited each detailed physician 9.79 (6.90) times per year in 2002-2004 to discuss Lipitor (Crestor), at an estimated cost of $150 ($187) per visit in 2003 dollars. This implies a "cost of relationship" of about $1,780.69 ($1,563.65) per physician-year in 2011 dollars, before accounting for the $50-$80 direct cost of payments.

Figure 5: Welfare Analysis with Other Frictions

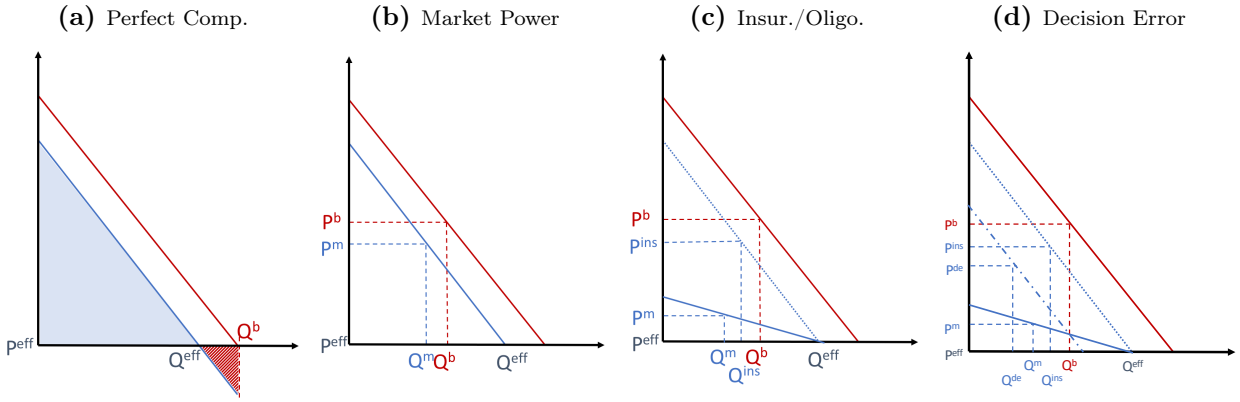

of this model is presented in panel (b) of Figure 5: a branded pharmaceutical manufacturer faces the residual demand curve in blue, which is again shifted outward in the presence of physician-firm payments. Market power causes "unbiased" quantities Q^m to be too low; thus, payments may increase prescribing toward the optimum $Q^m < Q^b < Q^{eff}$ (pictured) or cause prescribers to overshoot the optimum $Q^m < Q^{eff} < Q^b$. In the former case, the overall welfare impact of payments is positive, though consumer surplus declines; in the latter case, both total and consumer surplus decline.

Finally, we must also account for reasons that the "effective" demand curve for a given drug may not represent the appropriate one for welfare analysis. A leading example is insurance, pictured in panel (c) of Figure 5. The "true" demand curve is the solid blue line; the insured residual demand curve is the dotted blue line (which is significantly less elastic with respect to the point-of-sale price, as insurance enrollees bear only a fraction of that price out of pocket); and the "biased" demand curve is again in red. In this hypothetical, payments from firms reinforce the effects of insurance, each increasing consumption above the uninsured equilibrium: $Q^m < Q^{ins} < Q^b$.[32] The welfare implications are again ambiguous, and the consumer surplus effects of firm payments will depend on pass-through of producer prices to enrollee premiums.

In our supply analysis and counterfactuals, we account for the details of patient insurance and strategic interaction, and model point-of-sale prices as determined via bilateral bargaining between insurance plans and differentiated pharmaceutical suppliers. Point-of-sale prices then pass through partially into consumer out-of-pocket costs via a fixed cost-sharing rate.

[32] Another relevant extension would include the effect of strategic behavior of competitors. For example, in oligopoly, the residual demand curve can be distorted due to competitor pricing or payment behavior. This is the phenomenon highlighted in Inderst and Ottaviani (2012), where payments may even increase consumer surplus by improving allocative efficiency.

In this way, the basic machinery of our supply and demand model accounts for several economic forces that may cause inefficient utilization in equilibrium even absent payments to physicians.

The general point of panel (c) also extends beyond insurance, though. A large literature in economics and health services research has documented that healthcare decisions can be biased relative to the patient's optimum, due to a variety of potential frictions. These include physician information and skill (Abaluck et al. 2016; Chan Jr. et al. 2019; Currie and MacLeod 2020); imperfect agency (beyond the impact of detailing/meals); and "behavioral" errors such as present bias, symptom salience, and false beliefs (see Baicker et al. (2015) for a review). These biases could be positive or negative, depending on the context. In the case of statins, there is evidence of likely underprescribing relative to the clinical optimum (American College of Cardiology 2017). Motivated by this, panel (d) of Figure 5 shows one hypothetical extension of panel (c), grouping these "other" frictions under the term "Decision Error" for the sake of brevity and convenience. In this example, a negative decision error causes quantity to be too low absent payments, and payments increase quantity toward efficient levels, such that $Q^{de} < Q^{b} < Q^{eff}$. In our welfare analysis, we add a "decision error" parameter that allows for a range of assumptions on how payments might counteract, overshoot, or reinforce any baseline biases. We also show how to combine estimates from the clinical literature with our revealed preference demand estimates to inform that parameter.

5.2 Consumer Surplus

In order to derive consumer surplus implications of the demand model outlined and estimated in Section 4, we want to take seriously the many potential ways in which decision errors might drive a wedge between decision utility describing the combined physician/patient choice function and realized, welfare-relevant utility. We also want to consider how meals might counteract or reinforce such errors. To do so, we generalize our mean utility specification from Equation 1. We allow for the demand unobservable to have two components:

$$\xi_{djt} = \tilde{\xi}_{djt} + \varepsilon^{de}_{djt} \tag{7}$$

where $\tilde{\xi}_{djt}$ is a typical demand unobservable that impacts both choices and true realized utility, but ε^{de}_{djt} is a "decision error" in the spirit of Baicker et al. (2015) that impacts choice utility but not realized utility. That is, ε^{de}_{djt} affects consumer decisions, but it does not affect consumer surplus directly.

Given this model, we represent expected consumer surplus as:

$$CS_{dt}(\mathcal{J}_t) = \underbrace{Q_{dt}\frac{1}{\theta^p}\ln\left(1+\left(\sum_{j\in\mathcal{J}_{dt}} e^{\frac{\delta_{djt}}{1-\lambda}}\right)^{1-\lambda}\right)}_{CS \text{ implied by decision utility}} - \underbrace{\sum_{j\in\mathcal{J}_{dt}} q_{djt}\left(\frac{\varepsilon^{de}_{djt}+\theta^m_{dj}1_{\{m_{dj}>0\}}}{\theta^p}\right)}_{\text{adjustment for "decision errors" and meals}} \quad . \quad (8)$$

This is the standard formula derived by McFadden (1978), with a modification that captures the extent to which any meal payment effect causes prescribing to be more (vs. less) appropriate, conditional on all other variables. The first term reflects the consumer surplus that would be implied by our demand estimates if decision utility were equivalent to actual utility. The second term adjusts consumer surplus for the presence of a decision error that results in under- ($\varepsilon^{de}_{djt} < 0$) or over-prescribing ($\varepsilon^{de}_{djt} > 0$) as in Figure 5(d) above, as well as the countervailing (or reinforcing) effect of meals.[33]

5.2.1 Specifying and Estimating Decision Errors ε^{de}_{djt}

The decision error parameter approach has some appealing features. It can capture many theoretical frictions in a reduced form way (Baicker et al. 2015; Mullainathan et al. 2012). It is empirically flexible in that one can estimate decision utility following typical revealed preference-based procedures and then consider how different types of decision errors affect welfare. In prior studies with decision errors, data on *unbiased* decision-makers are leveraged to estimate true equilibrium welfare for the whole sample (Allcott and Taubinsky 2015; Bronnenberg et al. 2015; Handel and Kolstad 2015). We discuss how outside data might be used to calibrate a decision error in (the many) empirical contexts such as ours where no unbiased decision-makers are identified.

In our context, the important dimensions of the decision error specification are the mean decision error, heterogeneity in errors across physicians and molecules, and the correlation with meal payment effects. For example, $\varepsilon^{de}_{djt} = 0$ would be a case with no decision error at all, where meals simply bias utilization of promoted drugs upward. By contrast, $\varepsilon^{de}_{djt} = -\theta^m_{dj}$ would be a case where meals perfectly correct prescribing errors among those who receive them. To provide further intuition, Appendix C.3 presents a simple graphical model of the welfare consequences of meal effects, given physicians' underlying tendencies to over- or under-prescribe a given drug.

We study the welfare implications of two different specifications of decision errors. In our

[33] A related (and not mutually exclusive) interpretation would be that physicians maximize a sum of physician (chooser) and patient (consumer) utility, with ε^{de}_{dj} governing the difference between the physician's maximum and and the patient's.

main specification, we set a constant decision error across all doctors and statins $\varepsilon^{de}_{djt} := \varepsilon^{de}$, and we simulate counterfactuals for a range of decision errors, from substantial underprescribing to overprescribing.[34] This specification, while simple, has the virtue of being easy to interpret, and accommodates the finding in the prior literature that statins *as a drug class* are underprescribed (Baicker et al. 2015).

In an alternative specification, we simulate welfare under the assumption that decision errors are a scaled function of estimated physician-specific meal responses $\varepsilon^{de}_{d} := \gamma^{de}\bar{\theta}^{m}_{d}$, varying scalar γ^{de} to again allow for a range of potential under- or over-prescribing in the absence of meals. In this specification, $\gamma^{de} = -1$ represents a special case where meal payments perfectly correct for a given physician's average tendency to under/over prescribe, for those drugs for which meals are received. This specification also has cases where meals distort prescribing away from otherwise optimal behavior $\gamma^{de} = 0$, undercorrect $\gamma^{de} < -1$ or overcorrect $\gamma^{de} \in (-1, 0)$ underprescribing, and so on. Ultimately, the two types of specifications we explore do not result in different qualitative takeaways regarding the overall welfare effects of meals, so we do not explore other potential decision error specifications.

In both models of decision errors we study, we also compute the mean level of decision error that calibrates the expected total surplus a 30-day supply of statins generates in the model to the value of such a supply implied by the medical literature. We consider this the best outside estimate of the mean level of decision errors in statin prescribing in our sample.

5.3 Pricing Pharmaceuticals

Considering a counterfactual ban on meal payments requires understanding the potential effects of such a ban on prices, which we address using a model of how prices are set in equilibrium. While the details of pharmaceutical supply chains are notoriously complicated, and modeling such complications is the focus of ongoing research, we seek to abstract from the details that seem less relevant for our purposes while capturing enough of the key economics of pharmaceutical pricing to generate credible estimates of the direction and magnitude of equilibrium price changes with a meal payment ban. To accomplish this, we develop a model of a supplier (an entity subsuming manufacturers, wholesalers, and pharmacies) negotiating with a buyer (subsuming pharmacy benefit managers (PBMs) and insurers). This captures the primary forces relevant to our research question, abstracting from some of the details of the upstream interactions between suppliers, and from insurer competition and insurance

[34] This is similar to the approach in Handel (2013), which simulates counterfactual welfare over a range of assumptions regarding whether the friction underlying an inertial demand response represents a true social cost. An interesting feature of our specification is that the decision error need not be correlated with or bounded by the estimated friction.

plan structure.[35]

Let the supplier's profit be: $\pi(p^{pos}_{jrt}) = \sum_{d\in r} q_{djt}(p^{pos}_{jrt}(1-\tau_{jt}) - mc_{jt})$, where τ_{jt} is the manufacturer rebate and mc_{jt} captures the cost of manufacturing and distributing the marginal unit of molecule j. p^{pos}_{jrt} is the point-of-sale price insurers pay for the drug, which we model as constant across providers within region r. We link the negotiated point-of-sale price and out-of-pocket price paid by enrollees via $p^{oop}_{djt} = cs_{djt}p^{pos}_{jrt}$, where cs_{djt} is a cost-sharing parameter that varies across markets and years, depending on product mix and insurer mix (discussed in detail in Appendix A.2). Note that cost-sharing is applied to POS prices *before* rebates are taken out. We assume that cs_{djt} is exogenous, and we hold it fixed in counterfactual analyses. We take the region r over which point-of-sale prices are negotiated to be the state. We do not observe the mix of Part D plans covering a given physicians' enrollees, but this level of geography accounts for price variation driven by the entry and pricing decisions of standalone Part D plans and Medicare Advantage plans.[36]

We assume that prices of substitute drugs in the market are determined in a simultaneous Nash Equilibrium of Nash Bargaining between suppliers and buyers. In the model, each price maximizes the Nash Product of the gains from trade for each supplier and buyer pair, taking other prices as given. The first-order conditions of this model (see Appendix C for details) generate pricing equations that can be represented by:

$$p^{pos}_{jrt}(1-\tau_{jt}) = mc_j + b_{jrt}\left[\left(1+\sum_{d\in r}\frac{\partial q_{djt}}{\partial p^{oop}_{djt}}\frac{p^{oop}_{djt}-mc_j}{\sum_{d\in r} q_{djt}}\right)\frac{\sum_{d\in r}\widetilde{CS}_{dt}(\mathcal{J}_{dt}) - \widetilde{CS}_{dt}(\mathcal{J}_{dt}\setminus j)}{\sum_{d\in r} q_{djt}} + p^{pos}_{jrt}(1-\tau_{jt}) - mc_j\right] \tag{9}$$

Here the term $b_{jrt} \in [0,1]$ is a bargaining ability parameter, weighting the extent to which the optimal price depends on supplier profits ($b_{jrt} = 1$) vs. the expected additional buyer surplus ($b_{jrt} = 0$) in the case that a contract is agreed to for product j: $\widetilde{CS}_{dt}(\mathcal{J}_{dt}) - \widetilde{CS}_{dt}(\mathcal{J}_{dt}\setminus j)$. Notice that quantities and thus elasticities are driven by physician/enrollee decision-making based on out-of-pocket price under insurance coverage p^{oop}, but the insurer and supplier negotiate over point of sale price p^{pos}.

We follow recent papers on insurer-hospital bargaining (e.g., Gowrisankaran et al. 2015)

[35] As discussed by Starc and Swanson (2020), both pharmacies and pharmaceutical manufacturers have market power, but relative market power of different suppliers varies by drug. Pharmacies make larger margins on generic drugs than on branded drugs, while branded manufacturers command higher markups (even net of rebates) than generic manufacturers. These details are captured in a reduced form sense by the bargaining and cost-sharing parameters. Our counterfactual analysis will hold these fixed. This assumes that banning meals to physicians does not change the fundamental supply chain of the pharmaceutical industry or the general treatment of branded and generic therapies in insurance plan formularies.

[36] Standalone Part D plans enter and negotiate prices in one of 34 Part D pricing regions; PDP regions are either single states or supersets of states. Medicare Advantage plans enter at the county level. States strike a balance between these two levels of aggregation.

by using a parameter $\alpha^{cs} \in [0,1]$ to capture the relative weight insurers place on consumer surplus, and subtracting plan costs from weighted consumer surplus.[37] We also include a parameter $\alpha^{de} \in [0,1]$ that allows for a range of assumptions regarding whether insurers are naive vs. sophisticated regarding the impact of decision errors and meals on true consumer surplus and that perceived by insurers:

$$\begin{aligned}\widetilde{CS}_{dt}(\mathcal{J}_{dt}) \;:=\; & \alpha^{cs}\left[Q_{dt}\frac{1}{\theta^p}\ln\left(1+\left(\sum_{j\in\mathcal{J}_{dt}} e^{\frac{\delta_{djt}}{1-\lambda}}\right)^{1-\lambda}\right) - \alpha^{de}\sum_{j\in\mathcal{J}_{dt}} q_{djt}\left(\frac{\varepsilon^{de}_{dj}+\theta^m_{dj}1_{\{m_{dj}>0\}}}{\theta^p}\right)\right] \\ & -\sum_j q_{djt}(p^{pos}_{jrt}-p^{oop}_{djt})\ .\end{aligned}$$

This model accommodates the fully "naive" case where insurers negotiate prices under the assumption doctors know best ($\alpha^{de}=0$), as well as the fully "sophisticated" case where insurers perfectly adjust consumer surplus for decision errors and meals ($\alpha^{de}=1$), and every case in between.

5.3.1 Supply Model Estimation

The demand model estimates provide the utility parameters needed to compute demand elasticities and consumer surplus in the equilibrium observed in the data. They can also be used to estimate market shares and consumer surplus under counterfactual scenarios where any given product j is removed from the choice set, but prices of the remaining products stay the same. These are the critical inputs needed for the bargaining model. The remaining terms in the supply model are the bargaining ability weights (b_{jrt}), the insurer concern for consumer surplus vs. profits (α^{cs}), the decision error (ε^{de}), the pricing sophistication parameter (α^{de}), the manufacturer rebates (τ_{jt}), and the marginal costs (mc_{jt}).

To estimate the model for a given vector $(\varepsilon^{de}, \alpha^{de}, \tau_{jt}, mc_{jt})$, we parameterize bargaining ability parameters as a function of product and regional fixed effects, and specify the econometric unobservable as the residual variation in bargaining parameters needed to fit the model to the data. We then estimate the insurer concern for consumer surplus and bargaining ability parameters via GMM, using consumer surplus measures calculated at average prices in other regions as instruments to avoid potential simultaneity bias.[38] In Sections 5.4 and 5.5, we show welfare for a large range of potential values of ε^{de}. In the main text, we

[37] In contrast to these papers, we model pricing of a single product class (statins) rather than a bundle of products. Thus α^{cs} in our setting may also capture how plan enrollment might respond to disagreement in this particular product class.

[38] The fact that consumer surplus is a function of price can create an endogeneity problem where the surplus measures are correlated with the unobservable in the supply pricing equation.

assume $\alpha^{de} = 0$. Appendix H presents results with $\alpha^{de} = 1$, which is an interesting model in that an informed insurer can counteract some biases in prescribing through prices.

Unfortunately, we have neither the data nor the degrees of freedom to estimate (τ_{jt}, mc_{jt}) while simultaneously estimating a flexible set of bargaining weights (b_{jrt}). Unobserved rebates are an endemic challenge to research on pharmaceutical pricing, and the empirical difficulty of separately identifying bargaining weights and marginal costs is well-known (Gowrisankaran et al. 2015; Grennan 2013). Our solution is to use estimates of rebates and marginal costs from recent research on pharmaceutical markets, and we test sensitivity of our results to alternative assumptions. For example, in our baseline analysis, we assume that rebates for branded drugs were 26.3 percent, consistent with the average rebates for cardiovascular drugs reported to CMS in 2014, and we increase rebates to 48.3 percent for Lipitor in 2012 based on the estimates of post-patent expiration rebate increases in Arcidiacono et al. (2013). See Appendix E for full details. In our baseline analysis, we also assume that marginal costs for all jt are equal to 17 percent of the average POS price of generic statins: $mc = 0.17 * \overline{p^{pos}_{gen}}$. The value of 17 percent is taken from the average production costs of generic drugs in Sood et al. (2017), together with our assumption that the cost of producing a statin is invariant across molecules and branded/generic status. In Appendix H, we test robustness to reasonable alternative assumptions regarding (τ_{jt}, mc_{jt}) and our results are qualitatively unchanged.

5.3.2 Supply Estimation Results

Table 5 summarizes supply side parameter estimates. The most striking feature is the high bargaining parameter estimates for the branded drugs relative to generics. Because the generic sales are aggregated over firms, the bargaining parameters also capture within-molecule competitiveness. This can also be seen in the slightly larger bargaining parameter for generic atorvastatin, where only two manufacturers compete during the first six months of 2012, after which eleven more manufacturers enter. The larger bargaining parameters for Lipitor and Crestor in 2012 reflect the fact that POS prices remain high in many regions for much of 2012 as insurers are slow to adjust formularies, despite the improved outside option with generic atorvastatin entry.

Finally, we estimate that the weight insurers place on enrollee surplus in negotiations is equivalent to the weight they place on net costs: $\alpha^{cs} = 1$. This may reflect that enrollees are sensitive at the plan choice stage to formulary inclusion of important drugs. Indeed, Olssen and Demirer (2019) document substantial plan switching based on which statin brands are on formulary in Medicare Part D.

Table 5: Supply Parameter Estimates

		Atorvastatin	Lipitor	Lovastatin	Pravastatin	Crestor	Simvastatin
B_{2011}		–	0.60	0.08	0.07	0.60	0.06
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
B_{2012}		0.20	0.80	0.06	0.06	0.66	0.04
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
α^{cs}	1.00						
	(0.00)						

Notes: Based on 116,559 doctor-drug-brand-year observations with standard errors clustered at the doctor d level ($N_d = 13,793$) via delete-120 jackknife bootstrap.

5.4 Price and Quantity Effects of a Counterfactual Meal Ban

To better understand the economic effects of payments to physicians, we consider four counterfactual scenarios. The first scenario ("Ban, fix p") bans payments and computes new equilibrium quantities, but holds all prices fixed at those observed in the data. This allows us to isolate the effects of a ban on choice patterns alone. The second scenario ("Ban") allows point-of-sale and out-of-pocket prices and quantities to adjust to a new equilibrium. Comparing this to the observed data is the relevant exercise for considering the effects of a policy that bans meals—this comparison plays prominently in the next subsection, where we focus on the welfare implications of a ban. For the purpose of understanding the economics of a ban, comparing this to the first scenario helps to separate the equilibrium price response from the choice response. Our third and fourth scenarios set *out-of-pocket* prices equal to marginal costs and consider a ban and no ban with no price distortion ("Ban, $p = mc$" and "No Ban, $p = mc$", respectively). These scenarios provide a lens through which to consider the price distortions due to market power, holding payments fixed. They also provide approximations of an "efficient" benchmark—a payment ban and $p^{oop} = mc$ is efficient at one extreme where $\varepsilon^{de} \geq 0$, and no ban and $p^{oop} = mc$ is efficient if ε^{de} is negative and large enough. Table 6 displays several key results from these counterfactuals for 2011. 2012 results are qualitatively similar and shown in Appendix Table A11.

The top row of Table 6 shows total quantities of statins prescribed, highlighting several of the issues motivated in Section 5.1. A ban on meal payments reduces total statin usage as a share of cardiovascular claims by about one percentage point (whether or not we allow prices to adjust to a new equilibrium). This is a five percent reduction in total statin usage. For the focal branded statins, the increase is 29 percent on average. This is smaller than the large meal treatment effects documented in Section 4.4 above because not all doctors receive meals, and they only receive them for the branded statins. Some physicians also receive meals from both firms, and the business stealing effect mutes the overall effect of a meal ban

Table 6: Equilibrium Quantity and Price Effects of Meal Payments (2011)

	Observed	Ban, fix p	Ban	Ban, $p = mc$	No Ban, $p = mc$
$Q_{statins}$	0.182	0.173	0.173	0.185	0.197
	(0.000)	(0.002)	(0.002)	(0.002)	(0.000)
$Q_{Lipitor}$	0.040	0.033	0.033	0.044	0.053
	(0.000)	(0.001)	(0.001)	(0.002)	(0.000)
$Q_{Crestor}$	0.026	0.018	0.018	0.024	0.035
	(0.000)	(0.002)	(0.002)	(0.002)	(0.000)
$OOP_{statins}$	16.93	16.93	16.45	2.66	2.66
	(0.0)	(0.0)	(0.6)	(0.0)	(0.0)
$OOP_{Lipitor}$	32.35	32.35	31.51	2.66	2.66
	(0.1)	(0.1)	(0.8)	(0.0)	(0.0)
$OOP_{Crestor}$	31.10	31.10	29.92	2.66	2.66
	(0.0)	(0.0)	(0.18)	(0.0)	(0.0)
$POS_{statins}$	69.21	69.21	67.18	132.46	150.14
	(0.1)	(0.1)	(0.26)	(2.42)	(1.49)
$POS_{Lipitor}$	140.90	140.90	137.24	282.40	317.81
	(0.2)	(0.2)	(0.34)	(4.85)	(3.27)
$POS_{Crestor}$	133.75	133.75	128.70	265.0	305.0
	(0.1)	(0.1)	(0.77)	(5.80)	(3.12)

Notes: Authors' calculations of observed and counterfactual equilibrium statin quantities (share of cardiovascular utilization) and prices (point-of-sale and out-of-pocket price per 30-day supply), per supply and demand model described in text. 2011 only. All simulations assume $\alpha^{de} = 0$, implying that prices and quantities are invariant to ε^{de}. "Ban, fix p" eliminates meals, holding POS and OOP prices fixed. "Ban" eliminates meals and allows both prices and quantities to adjust. "Ban, $p = mc$" eliminates meals and sets p^{oop} equal to marginal costs, then allows POS prices and quantities to adjust. Finally, "No Ban, $p = mc$" simply sets p^{oop} equal to marginal costs, then allows POS prices and quantities to adjust. Estimates based on 116,559 doctor-drug-brand-year observations with standard errors clustered at the doctor d level ($N_d = 13,793$) via delete-120 jackknife bootstrap.

on branded drug usage in equilibrium.

The quantity estimates also show that pricing above marginal cost reduces total statin usage by about 1.5 percent of cardiovascular claims with meals (compare "Observed" to "No Ban, $p = mc$"). The price distortion is slightly smaller without meals (1.2 percentage points, compare "Ban" to "Ban, $p = mc$") because the price distortion is largest for the branded drugs, and there is less branded drug usage without meals. For both Lipitor and Crestor, meals counteract the fact that, even with cost sharing, insured patients face prices above marginal cost, resulting in total quantities that are closer to the efficient allocation. In the Lipitor case, meals cause utilization to undershoot the efficient allocation; in the Crestor case, meals cause utilization to fall between the allocations with $p = mc$ with and without meals. We explore the welfare implications of these changes in more detail in Section 5.5 below.

The other rows of Table 6 show the OOP prices faced by consumers and POS prices paid by insurers, respectively, under the counterfactual scenarios we study. They report quantity-weighted average prices across regions for all statins and also for Lipitor and Crestor independently. For both Lipitor and Crestor, banning meals causes point-of-sale prices to fall by \$4-\$5 per 30-day supply, which is about 3-4 percent in relative terms. The effect of a meal ban on out-of-pocket price is even smaller at around -\$1 per month supply, due to imperfect pass-through of point-of-sale prices into out-of-pocket consumer prices. Although meals shift the demand curve outward substantially, the effect of this demand expansion on price is dampened by the role of insurers as intermediaries negotiating point-of-sale prices. Physician/patient sensitivity to out-of-pocket still plays a role in suppliers' market power, however—if we counterfactually set $p = mc$ and divorce out-of-pocket prices from point-of-sale prices, point-of-sale prices for both branded drugs double.

It also bears noting that we observe different counterfactual pricing effects here, where we assume $\alpha^{de} = 0$ (i.e., insurers are "naive" regarding the roles of meals and decision errors in determining consumer surplus), vs. alternative specifications where we assume $\alpha^{de} = 1$ (i.e., insurers are "sophisticated" and prices are determined as a function of true consumer surplus). As noted above, Crestor's POS price goes down by \$5 under a meal ban regardless of ε^{de} when $\alpha^{de} = 0$. This is because, in negotiating prices, insurers consider decision utility to be equivalent to real utility, so a meal ban decreases the consumer surplus the insurer perceives to be associated with Crestor. As shown in Appendix Table A13, if insurers are instead sophisticated ($\alpha^{de} = 1$), then the effect of a meal ban is more nuanced. A ban leads to Crestor prices decreasing by \$2 if $\varepsilon^{de} = -350$ and *increasing* by \$2 if $\varepsilon^{de} = 0$. In the former case, meals offset underutilization of a high-value drug and the insurer internalizes that, paying higher prices in the presence of meals. In the latter case, meals just decrease

consumer surplus, and the insurer internalizes that and pays lower prices in the presence of meals.

5.5 Welfare Implications of a Counterfactual Meal Ban

To evaluate policies that seek to ban or limit meals and associated interactions perspective, we are interested in how price and quantity effects translate into welfare: consumer, producer, and total surplus. Because meals promote branded drugs, they are an expensive way to increase overall statin usage. Thus, the relationship between the above effects and consumer surplus depends crucially on the extent to which payment effects correct for decision errors that would otherwise lead to underutilization. In our baseline model, motivated by the American College of Cardiology's position that statins are underutilized overall (American College of Cardiology 2017), we suppose that all statins (and all physicians) are equally subject to the unidimensional decision error ε^{de} that biases substitution between statins and the outside option. This implicitly assumes that the quality rankings of different statins, conditional on out-of-pocket price, and the match value of particular statins to particular patient panels, are accurately reflected in our revealed preference estimates. In Appendix H, we alternatively assume a fixed ratio between ε_d^{de} and $\bar{\theta}_d^m$ (for example, physicians who are more responsive to meals may also be more subject to decision errors).

The surplus measures used in our welfare simulations are as follows. We present two different measures of consumer surplus. First, we show $CS_{dt}(\mathcal{J}_t)$ directly. Second, we subtract insurers' drug costs: $CS_{dt}(\mathcal{J}_t) - \sum_j q_{djt}(p_{jrt}^{pos} - p_{djt}^{oop})$. In this measure, when meals increase branded statin utilization and prices, the effect of that increase on drug expenditures fully enters Consumer Surplus (negatively). This second measure would be consistent with insurers' drug costs passing through fully to consumers (and/or the federal government, as Medicare Part D is a subsidized program) in the form of higher premiums.

We compute Producer Surplus as the total gains to producers from both OOP and POS payments, minus manufacturing, salesforce, and meal costs. As discussed above, we assume that marginal manufacturing costs are 17 percent of the average POS price of generic statins, and that salesforce costs are \$1,780.69 (\$1,563.65) per physician-year for Lipitor (Crestor) based on the estimates (discussed above) in Liu et al. (2020). As shown in Appendix H, Producer Surplus is higher under alternative cost assumptions, but Consumer Surplus is largely unchanged. Total Surplus is the sum of Producer Surplus and Consumer Surplus net of insurer drug costs. We calculate surplus for both 2011 and 2012. This highlights how the welfare effects of meal payments differ in the case of a single branded drug providing payments. It also provides some context for the magnitude of meal effects in that we can

compare them to the welfare impact of generic atorvastatin entry.

The results are summarized graphically in Figure 6 with all measures represented as percent changes relative to the baseline of the observed outcomes in 2011. Thus, "Observed 2011" is a flat line at zero, we compare "Observed 2012" to "Observed 2011" to quantify the welfare effect of atorvastatin entry, and we compare "Meal Ban t" to "Observed t" to quantify the welfare effect of a meal ban in year t. Appendix Table A12 provides estimates of Observed 2011 levels and standard errors on all estimates. The measures are represented in dollars per cardiovascular patient (many of whom will not receive a statin) in order to take into account changes on the extensive margin.

First, consider our measures of consumer surplus in panels (a) and (b). "Observed 2012" is rotated clockwise relative to "Observed 2011," reflecting that the benefit of atorvastatin entry in 2012 (and the associated price effects) will be decreasing in ε^{de}. Intuitively, the more negative ε^{de} is, the greater the implied benefit to patients of taking statins, and in turn of the statin market expansion in 2012. If $\varepsilon^{de} = 0$, Consumer Surplus increased by \$0.80 per patient due to generic atorvastatin entry, but Consumer Surplus net of transfers increased even more (\$2.86) due to the latter measure incorporating the full benefit of reduced POS prices.

In each year t, Consumer Surplus (with or without transfers) under a meal ban is rotated counter-clockwise, relative to Observed Consumer Surplus. Intuitively, more negative ε^{de} implies that statins are more valuable to patients, and hence that a meal ban has more potential to be harmful. The point at which the line for "Observed t" crosses the line for "Meal Ban t" is the point at which the benefits of increased statin use driven by meals, which disproportionately increase expensive branded statin use, exactly justifies the increased expenditures. Given the results of our supply and demand model, a ban improves consumer surplus in 2011 for $\varepsilon^{de}/\theta^p$ <-\$78, but only improves consumer surplus *net of transfers* for the more extreme threshold value of $\varepsilon^{de}/\theta^p$ <-\$270. In 2012, with a single branded firm offering meal payments, and another generic option available in atorvastatin, the necessary decision error is less dramatic, with cutoffs near -\$72 for pure Consumer Surplus and -\$202 for Consumer Surplus net of transfers. Intuitively, the large negative cutoffs of -\$270 and -\$202 reflect that, with our specification of a flat ε^{de} across all molecules, the *market expansion effect* of meals on overall statin use (which is much smaller than the effect of meals on use of *promoted statins*) must be valuable enough to justify increased utilization of expensive branded drugs.

From a producer surplus perspective (panel (c)), allowing meal payments is always preferred to a ban. We note that this is not a foregone conclusion, as business stealing effects can generate a prisoner's dilemma in which firms would prefer to ban advertising. The ef-

Figure 6: Welfare and Counterfactual Estimates

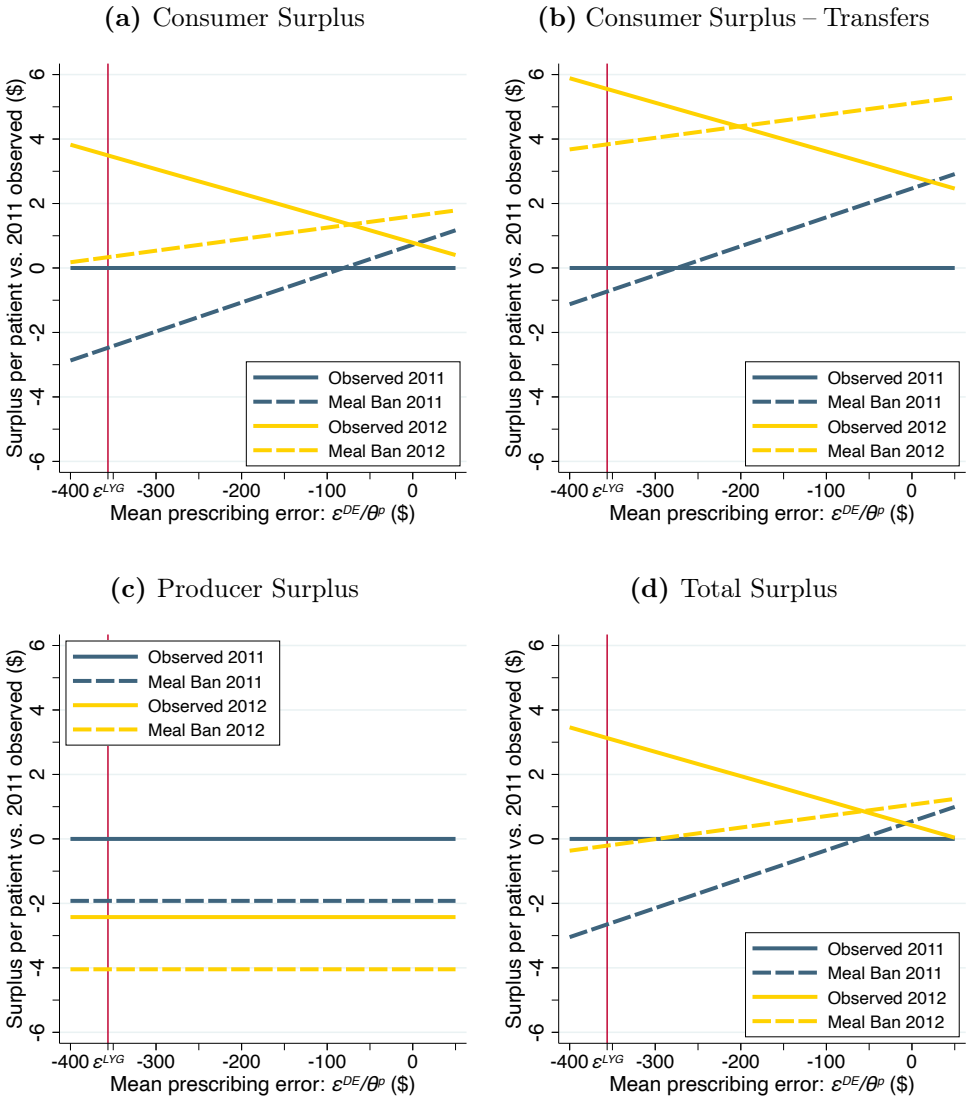

Notes: Authors' calculations of equilibrium surplus measures, in dollars per cardiovascular patient, relative to that Observed in 2011. "Meal Ban" counterfactuals allow both prices and quantities to adjust, per supply and demand model described in text. All simulations assume $\alpha^{de} = 0$; results shown for $\varepsilon^{de} \in [-\$400, \$50]$. Detailed results for select values of ε^{de} available in Appendix Table A12.

fect in the case we estimate here is fairly large, with a meal payment ban resulting in an approximately 19 percent decrease in producer surplus.

Taken together with Consumer Surplus net of transfers, panel (d) shows that, in the case of statins, meals increase Total Surplus as long as $\varepsilon^{de}/\theta^{p}$ <-\$58. If there is no underlying decision error, a meal ban increases Total Surplus by \$0.41 per patient. Alternatively, if meals cancel out decision errors on average among those receiving them ($\varepsilon^{de}/\theta^{p}$=-\$41), then the effect of meals encouraging statin use is almost exactly offset by the fact that meals encourage the use of expensive branded statins, resulting in a total surplus effect of meals that is economically small and statistically indistinguishable from zero.

Appendix Table A14 shows how our welfare simulations vary with our modeling assumptions, comparing the above ("Baseline") results to simulation results with alternative assumptions regarding rebates, marginal costs, $\alpha^{de} = 1$ instead of $\alpha^{de} = 0$ ("Pricing"), and $\varepsilon_d^{de} = \gamma^{de} * \bar{\theta}_d^m$ instead of fixed ε^{de} across all physicians. For each alternative specification, we show the effect of a meal ban on 2011 surplus (in dollars per cardiovascular patient) for a range of possible values of ε^{de}.[39] The effects of a meal ban are qualitatively and quantitatively similar across all modeling assumptions, with the exception that $\varepsilon_d^{de} = \gamma^{de} * \bar{\theta}_d^m$ implies that meals are are more likely to be welfare improving (e.g., Total Surplus increases for $\bar{\varepsilon}^{de} < -\$21$, vs. the lower threshold of -\$58 in our Baseline specification). This is not unexpected—if responsiveness to meals were highest among physicians with particularly large decision errors, then meals would be most effective where consumers stood to gain the most. Even here, though, the qualitative patterns remain the same.

5.5.1 Calibrating decision error magnitude using clinical data

The above counterfactual results make it clear that, for the statin market in 2011-12, meal payments from manufacturers to physicians increase demand for branded statins, and thus play an important role in generating profits for the manufacturers involved. They improve allocative efficiency by offsetting the distortion of high branded drug prices, but this is costly to consumers and insurers. Thus the effect of meal payments on consumer welfare depends on the extent to which meals simply increase usage vs. increase usage in cases where the drug would be severely underutilized $\varepsilon^{de} << 0$. Appendix C.3 shows this graphically. The extent of over- and underutilization (absent meals) surely varies across drugs. For statins, many studies point to potential underutilization, and this perspective is consistent with our result that meals tend to bring otherwise low prescribers closer to the prescribing behavior of those who do not receive meals. However, determining whether ε^{de} is sufficiently negative

[39] For the $\varepsilon_d^{de} = \gamma^{de} * \bar{\theta}_d^m$ specifications, the column value of ε^{de} is the average across sample physicians, given their average meal responsiveness $\bar{\theta}_d^m$.

for meals to be welfare-improving requires additional data.

We investigate this issue for our case study using estimates of the health benefits of statin regimens among indicated patients from clinical trials. The Heart Protection Study Collaborative Group indicates a benefit of a statin regimen of about 0.69 life years for Medicare-age enrollees if adherence is perfect over five years. Given conservative assumptions on adherence rates, benefits to non-adherent patients, and the dollar value of a life-year gain (see Appendix F for details), this implies that the decision-maker optimizing on behalf of an average indicated patient should compare the monthly out-of-pocket price of statins to a "flow" willingness-to-pay value of $516.[40] This measure suggests that the $\varepsilon^{de}/\theta^{p}$ consistent with statin demand in our data would be -$356 (indicated by a vertical red line in each panel of Figure 6). This is well below the 2011 threshold value of -$270 at which Consumer Surplus net of transfers improves in the presence of meals.

The estimate can be used directly in our framework, assuming that the Medicare cardiovascular patient population underlying our sample is similar to the population from which the life-year gain estimates are taken – UK adults over age 60, with blood total cholesterol concentrations of at least 135 mg/dL, and with coronary disease, other occlusive arterial disease, or diabetes (Heart Protection Study Collaborative Group 2009). We cannot provide direct evidence on this mapping using patient characteristics, but our simulations indicate that eliminating meals would reduce statin utilization by 5 percent, and the American College of Cardiology indicates that utilization of statins should *increase* by 24 percent from observed levels (American College of Cardiology 2017). That is, according to clinical guidelines, statin use is too low *even with* meals, and one might speculate that Medicare patients of cardiologists would be a natural population for the ACC's recommended expansion. We can also apply even more conservative assumptions to the mapping between the clinical data and the sample of Medicare cardiovascular patients we study. For example, Consumer Surplus net of transfers starts to decrease under a meal ban if more than 76 percent of (randomly selected) Medicare patients in our sample would experience the clinical benefits of statins from the medical literature, even if the remaining 24 percent experienced zero benefit. And so on.

We find this flexibility an appealing feature of the "decision error" approach to modeling frictions between decision utility and welfare relevant utility in health care. One can use a relatively transparent set of assumptions to map clinical data to revealed preference demand estimates for a given sample. For illustrative purposes, we have done so at the aggregate

[40] To obtain the "flow" value, we divide the total value of expected life-years gained from adherence over five years by the multiplier on monthly out-of-pocket costs that is necessary to cover five years of prescription statins, in present discounted value.

level and using one specific clinical trial result previously used in the economics literature (Sinkinson and Starc 2019). However, one could alternatively take a meta-analysis approach and show where the Consumer Surplus (or Total Surplus) threshold value falls relative to the distribution of clinical findings. One could also, perhaps with richer claims data, take further steps like matching patient observable characteristics in the prescribing data to those in the clinical trial data. Computing outcomes for a wide range of decision error values can also help explore robustness to assumptions.

6 Conclusion

In many industries, firms reach consumers through expert intermediaries. Interactions between firms and these experts, which can involve direct payments and other kinds of remuneration, risk creating conflicts of interest that can hinder efficiency. However, these interactions can also theoretically facilitate valuable information flows, reminders, or nudges, enhancing welfare. Further, they often take place in conjunction with other distortions due to agency, market power, and strategic interactions between firms. While recent theoretical work (Inderst and Ottaviani 2012) has shed new light on these tradeoffs, it has remained challenging to identify these relationships empirically, in part because of the strategic targeting of experts by firms. This gap in the literature is particularly important in light of recent policy debates over conflicts of interest in the U.S. health care and financial services industries.

We propose a strategy to overcome the challenges of empirically estimating these effects, and implement this strategy using an important case study in the health care industry. We introduce new instrumental variables, showing that local academic medical center conflict of interest policies influence the probability of payments from pharmaceutical companies for unaffiliated doctors in the same region. We employ machine learning methods to use this continuous instrumental variable to trace out the distribution of marginal treatment effects of firm payments to physicians in the market for statins. We also exploit variation in statin drug market structure over time, using the Lipitor patent expiration and ensuing generic entry to disentangle market power effects. Leveraging this approach with detailed data on prescriptions, prices, and payments, we are able to identify the impact of payments on prescribing behavior and welfare under a range of assumptions.

Overall, we find substantial heterogeneity across physicians in the expected response to payments, with the 90th percentile equivalent to roughly one standard deviation in prescribing heterogeneity and the 10th percentile not statistically different from zero. Firms target payments to physicians with: a larger expected response, larger patient panels, and lower than average prescribing of the focal drug. In short, firms successfully target payments to

physicians who will be responsive to their interactions and don't target those who aren't worth the expense.

Interestingly, these payments seem to mostly raise prescribing among targeted physicians such that they resemble those not targeted. This is at least consistent with arguments that payments are paired with information or reminders that might improve prescribing. To investigate this more precisely, we introduce a "decision error" parameter governing the extent to which payments interact with any baseline over- or underprescribing, and we compare welfare under the observed regime to a counterfactual regime with a payment ban. Payments improve allocation by offsetting the distortion of high prices for on-patent drugs. However, much of the gain accrues to manufacturers. Consumers only gain if payments offset underprescribing.

We then calibrate the decision error parameter to clinical data on the value of statins. This exercise suggests that, under reasonable assumptions about the mapping between patients in our prescribing data and participants in clinical trials, meal payments promoting branded statins increase consumer surplus as well as producer surplus. The magnitudes of these effects are similar to the introduction of generic atorvastatin.

There are limitations to our approach. We focus on a particular market, cardiologists and statin prescriptions, during a two-year time period near the expiration of the Lipitor patent. The dynamics of this market could differ in important ways from other drug and device markets in health care, and other industries where expert intermediaries play an important role, such as financial services. Future research can help to expand the scope of contexts studied and accumulate further policy-relevant evidence.

Can our results inform policy about banning meals and accompanying interactions more broadly? Of course any extrapolation should be done with caution, but we think that there are some more general lessons that can be learned. Our results suggest that a ban could harm consumer welfare in some markets. To evaluate a blanket ban, these harms would have to be balanced against the benefits of eliminating meals in markets with small, null, or even positive underlying decision errors. For example, there is evidence that Purdue's marketing of OxyContin to physicians had devastating effects on welfare, with repercussions that endure today (Alpert et al. 2019). Alternatively, perhaps policies that allow meal payments based on the state of clinical evidence relative to the current market uptake would remove the need to balance harms across markets using blanket policy. Of course, such policies would be much more difficult to administer. This idea is broadly consistent, though, with policies at some AMCs that try to encourage certain types of interactions and information exchange between industry and physicians.

Much can be gained from future research looking at similar phenomena in different con-

texts. The ability of pharmaceutical sales to target physicians seems extremely important. Given the ubiquitous findings of heterogeneity in treatment patterns across areas of medicine, this phenomenon may also extend beyond just pharmaceuticals. The spillovers identification strategy used here is fairly general, suggesting it could be used in many cases. As data on payments and treatment at finer timing units becomes available, future research may even be able to more clearly understand some of the dynamics that underlie these processes.

We also find the approach of calibrating revealed preference estimates to clinical data a potentially promising one for health care research. It is relatively straightforward, clear, and simple to implement in the manner we have done here. With increasingly rich clinical and real world treatment data becoming available in health care more broadly, this may offer one way to model welfare in the presence of concerns about various frictions and potential errors in patient care decisions.

References

Abaluck, J., Agha, L., Kabrhel, C., Raja, A., and Venkatesh, A. (2016). The determinants of productivity in medical testing: Intensity and allocation of care. *American Economic Review*, 106(12):3730–3764.

Abaluck, J., Bravo, M. M. C., Hull, P., and Starc, A. (2020). Mortality effects and choice across private health insurance plans. NBER Working Paper No. 27578.

Abaluck, J., Gruber, J., and Swanson, A. (2018). Prescription drug use under Medicare Part D: A linear model of nonlinear budget sets. *Journal of Public Economics*, 164:106–138.

Agha, L. and Zeltzer, D. (2019). Drug diffusion through peer networks: The influence of industry payments. *NBER Working Paper*, No. 26338.

Aitken, M. L., Berndt, E. R., Bosworth, B., Cockburn, I. M., Frank, R., Kleinrock, M., and Shapiro, B. T. (2018). The regulation of prescription drug competition and market responses. *Measuring and Modeling Health Care Costs*, 76:243.

Allcott, H. and Taubinsky, D. (2015). Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market. *American Economic Review*, 8(2501-2538).

Alpert, A., Evans, W. N., Lieber, E. M. J., and Powell, D. (2019). Origins of the opioid crisis and its enduring impacts. *NBER Working Paper*, No. 26500.

Altonji, J. G., Elder, T. E., and Taber, C. R. (2005). An evaluation of instrumental variable strategies for estimating the effects of catholic schooling. *Journal of Human Resources*, 40:791–821.

American College of Cardiology (2017). Study looks at uspstf vs. acc/aha guidelines for statin therapy. http://www.acc.org/latest-in-cardiology/articles/2017/04/18/15/45/study-looks-at-uspstf-vs-acc-aha-guidelines-for-statin-therapy.

Anagol, S., Cole, S., and Sarkar, S. (2017). Understanding the advice of commissions-motivated agents: Evidence from the indian life insurance market. *Review of Economics and Statistics*, 99(1):1–15.

Andresen, M. E. (2018). Exploring marginal treatment effects: Flexible estimation using stata. *The Stata Journal*, 18(1):118–158.

Angrist, J., Lavy, V., and Schlosser, A. (2010). Multiple experiments for the causal link between the quantity and quality of children. *Journal of Labor Economics*, 28:773–824.

Arcidiacono, P., Ellickson, P. B., Landry, P., and Ridley, D. B. (2013). Pharmaceutical followers. *International Journal of Industrial Organization*, 31(5):538–553.

ASPE (2016). Observations on trends in prescription drug spending. Available at https://aspe.hhs.gov/pdf-report/observations-trends-prescription-drug-spending.

Baicker, K., Mullainathan, S., and Schwartzstein, J. (2015). Behavioral hazard in health insurance. *Quarterly Journal of Economics*, 130(4):1623–1667.

Belloni, A., Chernozhukov, V., Fernandez-Val, I., and Hansen, C. (2017). Program evaluation and causal inference with high-dimensional data. *Econometrica*, 85(1):233–298.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 25(2):242–262.

Berry, S. T. and Waldfogel, J. (1999). Public radio in the united states: Does it correct market failure or cannibalize commercial stations? *Journal of Public Economics*, 71(2):189–211.

Bound, J. and Jaeger, D. A. (2000). *Do compulsory school attendance laws alone explain the association between quarter of birth and earnings?*, volume 19 of *Research in Labor Economics*, pages 83–108. Bingley, West-Yorkshire: Emerald Group Publishing Limited.

Brinch, C. N., Mogstad, M., and Wiswall, M. (2017). Beyond LATE with a discrete instrument. *Journal of Political Economy*, 125(4):985–1039.

Bronnenberg, B. J., Dubé, J.-P., Gentzkow, M., and Shapiro, J. M. (2015). Do pharmacists buy bayer? informed shoppers and the brand premium. *Quarterly Journal of Economics*, 130(4):1669–1726.

Campbell, J. J. (2008). *Understanding Pharma*. Pharmaceutical Institute, Raleigh, NC, 2nd edition.

Cardell, N. S. (1997). Variance components structures for the extreme-value and logistic distributions with application to models of heterogeneity. *Econometric Theory*, 13(2):185–213.

Carey, C. M., Lieber, E. M. J., and Miller, S. (2020). Drug firms' payments and physicians' prescribing behavior in Medicare Part D. *NBER Working Paper*, No. 26751.

Chan Jr., D. C., Gentzkow, M., and Yu, C. (2019). Selection with variation in diagnostic skill: Evidence from radiologists. NBER Working Paper No. 26467.

Chandra, A., Finkelstein, A., Sacarny, A., and Syverson, C. (2016). Health care exceptionalism? performance and allocation in the us health care sector. *American Economic Review*, 106(8):2110–44.

Chandra, A., Gruber, J., and McKnight, R. (2010). Patient cost-sharing and hospitalization offsets in the elderly. *American Economic Review*, 100(1):193–213.

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *Econometrics Journal*, 21(1):C1–C68.

Chintagunta, P. and Manchanda, P. M. (2004). Responsiveness of physician prescription behavior to salesforce effort: An individual level analysis. *Marketing Letters*, 15(2-3):129–145.

Clemens, J. and Gottlieb, J. D. (2014). Do physicians' financial incentives affect medical treatment and patient health? *American Economic Review*, 104(4):1320–1349.

Colantonio, L. D., Huang, L., Monda, K. L., and et al. (2017). Adherence to high-intensity statins following a myocardial infarction hospitalization among medicare beneficiaries. *JAMA Cardiology*, 2(8).

Colantonio, L. D., Rosenson, R. S., Deng, L., Monda, K. L., Dai, Y., Farkouh, M. E., Safford, M. M., Philip, K., Mues, K. E., and Muntner, P. (2019). Adherence to statin therapy among us adults between 2007 and 2014. *Journal of the American Heart Association*, 8(1).

ConsumerReports (2014). Evaluating statin drugs to treat: High cholesterol and heart disease.

Currie, J. M. and MacLeod, W. B. (2020). Understanding doctor decision making: The case of depression treatment. *Econometrica*, 88(3):847–878.

Cutler, D. M. (2004). *Your Money or Your Life*. Oxford University Press.

Cutler, D. M., Skinner, J. S., Stern, A. D., and Wennberg, D. (2019). Physician beliefs and patient preferences: A new look at regional variation in health care spending. *American Economic Journal: Economic Policy*, 11(1):192–221.

Datta, A. and Dave, D. (2016). Effects of physician-directed pharmaceutical promotion on prescription behaviors: Longitudinal evidence. *Health Economics*.

David, G., Markowitz, S., and Richards-Shubik, S. (2010). The effects of pharmaceutical marketing and promotion on adverse drug events and regulation. *American Economic Journal: Economic Policy*, 2(4):1–25.

Deichmann, R. E., Morledge, M. D., Ulep, R., Shaffer, J. P., Davies, P., and van Driel, M. L. (2006). A metaanalysis of interventions to improve adherence to lipid-lowering medication. *Ochsner Journal*, 16(3):230–237.

DeJong, C., Aguilar, T., Tseng, C.-W., Lin, G. A., Boscardin, W. J., and Dudley, R. A. (2016). Pharmaceutical industry-sponsored meals and physician prescribing patterns for medicare beneficiaries. *JAMA Internal Medicine*, 176(8):1114–1122.

Dickstein, M. J. (2016). Physician vs. patient incentives in prescription drug choice. *Mimeo*.

Dubois, P., Griffith, R., and O'Connell, M. (2018). The effects of banning advertising in junk food markets. *Review of Economic Studies*, 85:396–436.

Duggan, M. and Scott Morton, F. (2010). The impact of medicare part d on pharmaceutical prices and utilization. *American Economic Review*, 100(1):590–607.

Egan, M., Ge, S., and Tang, J. (2020). Conflicting interests and the effect of fiduciary duty – evidence from variable annuities. NBER Working Paper No. 27577.

Egan, M., Matvos, G., and Seru, A. (2019). The market for financial adviser misconduct. *Journal of Political Economy*, 127(1):233–295.

Einav, L., Finkelstein, A., and Polyakova, M. (2018). Private provision of social insurance: Drug-specific price elasticities and cost sharing in medicare Part D. *American Economic Journal: Economic Policy*, 10(3).

Fugh-Berman, A. and Ahari, S. (2007). Following the script: How drug reps make friends and influence doctors. *PLoS Med*, 4(4):e150.

Goldman, D., Joyce, G., and Zheng, Y. (2007). Prescription drug cost sharing. *JAMA*, 298(1):61–69.

Gowrisankaran, G., Nevo, A., and Town, R. (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review*, 175(1):172–203.

Grande, D., Frosch, D. L., and Perkins, A. W. (2009). Effect of exposure to small pharmaceutical promotional items on treatment preferences. *JAMA Internal Medicine*, 169(9):887–893.

Grassley, C. (2009). Grassley works for disclosure of drug company payments to medical groups. https://www.grassley.senate.gov/news/news-releases/grassley-works-disclosure-drug-company-payments-medical-groups.

Grennan, M. (2013). Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review*, 103(1):145–177.

Grennan, M. (2014). Bargaining ability and competitive advantage: Empirical evidence from medical devices. *Management Science*, 60(12):3011–3025.

Gruber, J. and Owings, M. (1996). Physician financial incentives and cesarean section delivery. *RAND Journal of Economics*, 27:99–123.

Guo, T., Sriram, S., and Manchanda, P. (2020). "Let the sunshine in": The impact of industry payment disclosure on physician prescription behavior. *Marketing Science*, 39(3):459–665.

Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, 103(7):2643–82.

Handel, B. R. and Kolstad, J. (2015). Health insurance for 'humans': Information frictions, plan choice, and consumer welfare. *American Economic Review*, 105(8):2449–2500.

Handel, B. R. and Schwartzstein, J. (2018). Frictions or mental gaps: What's behind the information we (don't) use and when do we care? *Journal of Economic Perspectives*, 32(1):155–178.

Hastings, J., Hortacsu, A., and Syverson, C. (2017). Sales force and competition in financial product markets: The case of mexico's social security privatization. *Econometrica*, 85(6):1723–1761.

Heart Protection Study Collaborative Group (2009). Statin cost-effectiveness in the united states for people at different vascular risk levels. *Circulation: Cardiovascular Quality and Outcomes*, 2.

Heckman, J. J., Urzua, S., and Vytlacil, E. (2006). Understanding instrumental variables in models with essential heterogeneity. *Review of Economics and Statistics*, 88(3):389–432.

Heckman, J. J. and Vytlacil, E. J. (2007). Econometric evaluation of social programs, part ii: Using the marginal treatment effect to organize alternative econometric estimators to evaluate social programs, and to forecast their effects in new environments. *Handbook of Econometrics*, 6:4875–5143.

Ho, K. and Pakes, A. (2014). Hospital choices, hospital prices, and financial incentives to physicians. *American Economic Review*, 104(12):3841–84.

Hoadley, J., Summer, L., Hargrave, E., Cubanski, J., and Neuman, T. (2014). Medicare part d in its ninth year: The 2014 marketplace and key trends, 2006-2014. Technical report, Kaiser Family Foundation.

Iizuka, T. (2012). Physician agency and adoption of generic pharmaceuticals. *American Economic Review*, 102(6):2826–58.

Inderst, R. and Ottaviani, M. (2012). Competition through commissions and kickbacks. *American Economic Review*, 102(2):780–809.

Jacobson, M., O'Malley, A. J., Earle, C. C., Gaccione, P., Pakes, J., and Newhouse, J. P. (2006). Does reimbursement influence chemotherapy treatment for cancer patients. *Health Affairs*, 25(2):437–443.

Johnson, E. M. and Rehavi, M. M. (2016). Physicians treating physicians: Information and incentives in childbirth. *American Economic Journal: Economic Policy*, 8(1):115–41.

Kakani, P., Chernew, M., and Chandra, A. (2020). Rebates in the pharmaceutical industry: Evidence from medicines sold in retail pharmacies in the U.S. *NBER Working Paper*, No. 26846.

Kesselheim, A. S., Avorn, J., and Sarpatwari, A. (2016). The high cost of prescription drugs in the united states. *JAMA*, 316(8):858–871.

Khedkar, P. and Sturgis, M. (2015). For pharma reps, a challenging market for physician access gets even tougher. Accessmonitor 2015 executive summary, ZS Associates.

King, M. and Bearman, P. S. (2017). Gifts and influence: Conflict of interest policies and prescribing of psychotropic medications in the united states. *Social Science and Medicine*, 172:153–162.

Kremer, S. T., Bijmolt, T. H., Leeflang, P. S., and Wieringa, J. E. (2008). Generalizations on the effectiveness of pharmaceutical promotional expenditures. *International Journal of Research in Marketing*, 25(4):234–246.

Larkin, I., Ang, D., and Steinhart, J. (2017). Association between academic medical center pharmaceutical detailing policies and physician prescribing. *JAMA*, 317(17):1785–1795.

Levitt, S. D. and Syverson, C. (2008). Market distortions when agents are better informed: The value of information in real estate transactions. *Review of Economics and Statistics*, 90(4):599–611.

Liu, Q., Gupta, S., Venkataraman, S., and Liu, H. (2020). An empirical model of drug detailing: Dynamic competition and policy implications. *Management Science*, 62(8):2321–2340.

Manchanda, P. M. and Honka, E. (2009). The effects and role of direct-to-physician marketing in the pharmaceutical industry: An integrative review. *Yale Journal of Health Policy, Law, and Ethics*, 5(2).

McFadden, D. (1978). Modeling the choice of residential location. *Transportation Research Record*, 673.

Midboe, A. M., Wu, J., Erhardt, T., Carmichael, J. M., Bounthavong, M., Christopher, M. L. D., and Gale, R. C. (2018). Academic detailing to improve opioid safety: Implementation lessons from a qualitative evaluation. *Pain Medicine*, 19:S46–S53.

Mizik, N. and Jacobson, R. (2004). Are physicians "easy marks"? quantifying the effects of detailing and sampling on new prescriptions. *Management Science*, 50(12):1704–1715.

Morgenson, G. (2020). It was his dream job. he never thought he'd be bribing doctors and wearing a wire for the feds. *NBCnews.com*.

Mullainathan, S., Schwartzstein, J., and Congdon, W. J. (2012). A reduced-form approach to behavioral public finance. *Annual Review of Economics*, 4:511–540.

Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica*, 46(1):69–85.

Narayanan, S. and Manchanda, P. (2009). Heterogeneous learning and the targeting of marketing communication for new products. *Marketing Science*, 28(3):424–441.

Olssen, A. L. and Demirer, M. (2019). Drug rebates and formulary design: Evidence from statins in Medicare Part D. *Mimeo*.

Ornstein, C. and Grochowski Jones, R. (2015). About the dollars for docs data.

Pew Charitable Trust (2013). Persuading the prescribers: Pharmaceutical industry marketing and its influence on physicians and patients. Projects: Pew Prescription Project.

Robles-Garcia, C. (2020). Competition and incentives in mortgage markets: The role of brokers. *Mimeo*.

Schneider, H. J. (2012). Agency problems and reputation in expert services: Evidence from auto repair. *Journal of Industrial Economics*, 60(3):406–433.

Scott Morton, F. and Kyle, M. (2012). Markets for pharmaceutical products. In McGuire, T., editor, *Handbook of Health Economics*, volume 2, chapter 12. Elsevier B.V.

Shapiro, B. (2018a). Informational shocks, off-label prescribing and the effects of physician detailing. *Management Science*, 64(12):5925–5945.

Shapiro, B. (2018b). Positive spillovers and free riding in advertising of prescription pharmaceuticals: The case of antidepressants. *Journal of Political Economy*, 126(1):381–437.

Sinkinson, M. and Starc, A. (2019). Ask your doctor? direct-to-consumer advertising of pharmaceuticals. *The Review of Economic Studies*, 86(2):836–881.

Sood, N., Shih, T., Nuys, K. V., and Goldman, D. (2017). The flow of money through the pharmaceutical distribution system. USC Schaeffer Center.

Spurling, G. K., Mansfield, P. R., Montgomery, B. D., Lexchin, J., Doust, J., Othman, N., and Vitry, A. I. (2010). Information from pharmaceutical companies and the quality, quantity, and cost of physicians' prescribing: a systematic review. *PLoS Med*, 7(10):e1000352.

Starc, A. and Swanson, A. (2020). Preferred pharmacy networks and drug costs. Conditionally accepted, American Economic Journal: Economic Policy.

The Health Strategies Consultancy LLC (2005). Follow the pill: Understanding the u.s. commercial pharmaceutical supply chain.

United States Department of Justice (2009). Eli lilly and company agrees to pay $1.415 billion to resolve allegations of off-label promotion of zyprexa. https://www.justice.gov/archive/opa/pr/2009/January/09-civ-038.html.

Wagner, M., Goetghebeur, M., Merikle, E., Pandya, A., Chu, P., and Taylor, D. C. (2009). Cost-effectiveness of intensive lipid lowering therapy with 80 mg of atorvastatin, versus 10 mg of atorvastatin, for secondary prevention of cardiovascular disease in canada. *Canadian Journal of Clinical Pharmacology*, 16(2):331–45.

Waldfogel, J. (2007). *The Tyranny of the Market*. Cambridge, MA: Harvard University Press.

Walter, M. (2020). A missed opportunity: Statin use low among patients with peripheral artery disease.

Yeh, J. S., Franklin, J. M., Avorn, J., Landon, J., and Kesselheim, A. S. (2016). Association of industry payments to physicians with the prescribing of brand-name statins in massachusetts. *JAMA Internal Medicine*, 176(6):763–768.

Yu, N. L., Atteberry, P., and Bach, P. B. (2018). Spending on prescription drugs in the US: Where does all the money go?

APPENDICES—FOR ONLINE PUBLICATION ONLY

A Additional Institutional Background

A.1 Medicare Part D

37 million people, or 70 percent of eligible Medicare beneficiaries, enrolled in Part D plans in 2014 (Hoadley et al. 2014). Medicare-eligible individuals can acquire prescription drug coverage through standalone Part D plans or bundled with medical and hospital coverage in the form of "Medicare Advantage" plans. Utilization of drugs in the Part D program is a function of physicians' prescribing decisions. These in turn may be impacted by: prescribers' training and knowledge, interactions with pharmaceutical firms, and preferences over cost control; the relevant drugs' effectiveness, side effects, and out-of-pocket costs; and Part D insurers' coverage policies.

Part D plans are offered by private insurers, but the federal Centers for Medicare and Medicaid Services mandates coverage generosity of plans in terms of actuarial value, types of drugs covered, and pharmacy network breadth. Enrollees are entitled to basic coverage of prescription drugs by a plan with equal or greater actuarial value to a standard Part D plan.[41]

The majority of Part D enrollees are not enrolled in standard plans, but rather in actuarially equivalent or "enhanced" plans with non-standard deductibles and tiered copays where enrollees' out-of-pocket costs vary across drugs and pharmacies. Branded drugs with close generic substitutes (e.g., Lipitor and Crestor vs. simvastatin and pravastatin prior to Lipitor's patent expiration) generally have higher copays than generics, while branded drugs with generic equivalents (e.g., Lipitor after patent expiration) have even higher copays or may not be covered by plans at all. On the other hand, approximately 30 percent of Part D enrollees qualify for low-income subsidies (LIS), which entitles them to substantial reductions in premiums and out-of-pocket costs on covered drugs; maximum copays for LIS enrollees are low or zero.[42]

Part D issuers receive premiums from enrollees and a variety of subsidy payments from

[41] In 2011, the standard plan covered: none of the first $310 in drug costs each year (the deductible); 75 percent of costs for the next $2,530 of drug spending (up to $2,840 total; the "initial coverage region"); 50 percent of branded costs for the next $3,607 of drug spending (up to $6,447 total; the "donut hole"); and 95 percent of costs above $6,447 in total drug spending (the "catastrophic region").

[42] Partial subsidies are available at 150 percent of the federal poverty level (FPL); full subsidies are available at 100 percent of FPL. LIS enrollees can enroll premium-free in "benchmark plans" or enroll in a non-benchmark plan and pay the difference between the chosen plan's premium and the benchmark premium out-of-pocket.

CMS: risk-adjusted direct subsidies for each enrollee, additional subsidies to cover LIS premiums and cost-sharing, and reinsurance for particularly high-cost enrollees. They also receive or pay "risk corridor" transfers such that the issuers' profits/losses are within certain bounds.[43] Although issuers' strategies and profits are heavily regulated by CMS, they can constrain costs through formulary design (drugs' coverage and placement on tiers, which determine patients' access to those drugs and out-of-pocket costs), negotiations with drug manufacturers, and negotiations with pharmacies.

A.2 Regional Prices and Formulary Variation in 2012

In our structural analyses, we identify the price sensitivity of demand using panel variation in out-of-pocket prices faced by Medicare enrollees. This variation is driven by Lipitor's patent expiration and by regional variation in insurers' responses to Lipitor's patent expiration.

Out-of-pocket prices are generally determined using insurance plan-specific formulas as a function of drug coverage, placement on tiers, point-of-sale price, and benefit phase. If a drug is covered, the unsubsidized out-of-pocket price will be *either* the tier-phase-specific copay *or* the product of the tier-phase-specific coinsurance and the point-of-sale price of the drug. Low-income subsidy enrollees face copay maximums as a function of their income. For example, in 2011, the maximum out-of-pocket price for LIS beneficiaries with income above 100 percent of the federal poverty level (FPL) was $2.50 for generic drugs and $6.30 for branded drugs, and many LIS beneficiaries qualified for more generous subsidies based on income.

For our model estimation, we use point-of-sale and out-of-pocket prices from the CMS Part D public use files for Q2 2011 and Q3 2012. In each file, we observe POS price for a 30-day supply, formulary tier placement, and unsubsidized beneficiary cost-sharing for each plan-drug, where drugs are identified by two codes: the RxNorm concept unique identifier (RXCUI) and the national drug code (NDC). These codes are quite specific. For example, NDC uniquely identifies the labeler (roughly, the pharmaceutical manufacturer); the specific strength, dosage form (i.e, capsule, tablet, liquid) and formulation, and the package size and type. We use the public use files to calculate out-of-pocket price per 30-day supply for an unsubsidized enrollee in each coverage region of the Medicare Part D plan benefit design, for each plan-year-drug code. For off-formulary drugs (i.e., drugs not covered by the plan at all), we set the out-of-pocket price equal to the point-of-sale price. To calculate the average unsubsidized (non-LIS) out-of-pocket price for each plan-drug-year, we feed the average spending for non-LIS enrollees in 2011 and 2012 from Starc and Swanson (2020) (Table 1)

[43] Insurers bear all upside/downside risk within a 5 percent band of zero profit; outside this risk corridor, the plan absorbs 20-25 percent of profits and losses.

through the nonlinear benefit structure in each plan-year to determine the weight to be put on each coverage phase-specific price. We limit LIS out-of-pocket prices to not exceed the maximum copays for branded and generic drugs (as appropriate) for non-institutionalized LIS beneficiaries with incomes over 100 percent of FPL.[44] Finally, we calculate an average out-of-pocket price per plan-drug-year by aggregating across non-LIS and LIS out-of-pocket prices, weighting by enrollment at the LIS status-plan-year level from the Medicare Advantage/Part D Contract and Enrollment Data.[45]

Given that our prescription drug claims data are at the prescriber level and thus cannot be linked to plans, we aggregate up to the state-drug-year level using plan enrollment data to construct weighted averages. Standalone Part D plans enter, negotiate prices, and set beneficiary cost-sharing in one of 34 Part D pricing regions; PDP regions are either single states or supersets of states. In contrast, Medicare Advantage plans enter at the county level. States strike a balance between these two levels of aggregation.

When Lipitor's patent expired in November 2011, generic atorvastatin was introduced by two generic manufacturers—the "authorized" generic firm Watson Pharmaceuticals and the paragraph IV challenger Ranbaxy Laboratories—that were afforded 180 days of exclusivity from other generic competition. After Lipitor's loss of exclusivity, essentially all Part D plans added atorvastatin to their formularies in 2012. Conversely, many plans did not immediately remove Lipitor from their formularies. In Q3 2012, 50 percent of plans still covered Lipitor. To the extent that some enrollees whose plans dropped Lipitor from their formularies were motivated to purchase Lipitor in cash (in which case the claim would not be recorded in the Medicare Part D data), this will bias our estimates of price sensitivity upward in magnitude. POS and OOP prices are summarized in Table 1 in the main text.

Cross-sectional variation in prices is generated by plan-pharmacy negotiations over point-of-sale prices and by plan-specific decisions regarding drug coverage and tiering. The coefficients of variation for the point-of-sale (out-of-pocket) price across Part D regions in 2011 was 0.02 (0.19) for Lipitor and 0.02 (0.16) for Crestor. The coefficients of variation for Lipitor and Crestor were similar in 2012. There was larger variation in 2012 in terms of both point-of-sale (CV =0.11) and out-of-pocket price (CV =0.24) for generic atorvastatin. This price variation, at the state-year-drug level, is presented for our focal drugs in Table A1 below.

Many of the determinants of both point-of-sale and out-of-pocket prices across regions at a point in time are likely driven by insurer-specific factors that are correlated across

[44] https://q1medicare.com/PartD-The-2014-Medicare-Part-D-Outlook.php

[45] https://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/MCRAdvPartDEnrolData

Table A1: Lipitor, Atorvastatin, and Crestor Prices—2011 to 2012

		2011				2012				2011–2012	
		mean	s.d.	$\beta^{\text{cross-sec.}}$	s.e.$^{\text{cross}}$	mean	s.d.	β^{cross}	s.e.$^{\text{cross}}$	β^{panel}	s.e.$^{\text{panel}}$
Lipitor	OOP	31.87	6.079	0.810	0.239	66.12	14.86	0.819	0.119	1.019	0.039
	POS	140.0	3.460	2.529	1.889	163.6	8.960	1.950	1.049	0.970	0.050
Atorva-	OOP					10.03	2.420	0.850	0.300		
-statin	POS					31.28	3.420	0.819	0.200		
Crestor	OOP	31.11	4.869	0.579	0.189	31.22	5.019	0.500	0.250	0.500	0.250
	POS	137.7	3.410	2.890	2.529	160.9	3.519	0.280	0.540	1.009	0.009

Notes: Reports state-year-drug out-of-pocket (OOP) and point-of-sale (POS) prices (means and standard deviations) and regressions of prices in one state (or state-year) on the prices of dominant insurers in other states, within-year ("cross") or across years within state ("panel").

regions. These might include management, contracts with prescription benefit managers, and costs. Given this, we introduce another source of identifying variation—for each plan-drug-state-year, we calculate the average price for that plan's issuer, for the same drug-year in *other* pricing regions, and we aggregate that instrument across plans within each state to generate a state-drug-year-specific instrument. The logic is as follows: if (for instance) United HealthCare were particularly slow to remove Lipitor from its formularies, then Lipitor prices in 2012 would be higher in regions dominated by United HealthCare for reasons unrelated to those regions' latent price-sensitivity or willingness to substitute to generic equivalents. The association between the point-of-sale and out-of-pocket prices within-year and across time within-state is in the "cross" and "panel" columns in Table A1 (β reports the "first stage" regression coefficient with the standard errors in the next column). There is a strong positive association between the pricing policies of the dominant insurers in each state and their pricing policies in other regions. This holds within each year, looking across states cross-sectionally ("cross"), and within states, looking across years, which we can see in the final "panel" column that pools years and controls for state fixed effects. These associations are generally more precise for OOP prices (which we use in our demand analysis) than for POS prices. This suggests that the correlation in "insurer-specific factors" across regions is stronger for benefit design (e.g., formulary structure) than for insurer-supplier POS price negotiations.

B Data Set—Construction and Context

B.1 From Full to Estimation Sample

Table A2 reports summary statistics for key prescribing and meal-payment variables. In terms of the two main regressions used to identify the demand parameters: the price and nest regression is based on data at the doctor-molecule-year level (djt; Panel a) for all drugs and uses the sample corresponding to column (3); the meal regression is based on data at the doctor-molecule level (dj; Panel b) only for Crestor and Lipitor and uses the sample corresponding to column (4).

B.2 Cleaning and Validating Payments Data

The payment data is based on publicly available data released by firms prior to the Sunshine Act-required reporting that began in 2013. When posting these reports, each firm adopted its own standards for specificity,[46] categorization approach,[47] and accuracy. Physician-level identifiers were ambiguous and often limited to a name, city of address, and perhaps a specialty. Furthermore, many of these documents have since been removed from easily accessible websites. During the period that these payments were still posted on the firms' websites, the enterprise software company Kyruus collected these reports as a part of their initiative to analyze physician-firm relationships. In order to create a disambiguated physician-level dataset using the unstandardized reports, Kyruus utilized their proprietary machine learning algorithms to match each individual-firm data point with the physician most likely to be the true recipient. The resulting dataset, generously provided to us by Kyruus, connects each physician-firm-payment to the most probable unique National Provider Identifier—a variable enabling us to merge this data to a number of other datasets.

There is significant heterogeneity in the nature of payments as they relate to the potential for conflict of interest. For example, a physician may receive a royalty payment for an invention sold by a company or a consulting payment for advice on product development. Other payments might not be related to a product at all. We construct two main categories of payments: "research" and "general" (all non-research payments). This scheme closely follows that of Open Payments and excludes all royalty payments. Within general payments we identify three sub-categories: "meals," "travel or lodging," and "consulting, speaking or education."

[46] For example, while many firms reported whole dollar amounts, Allergan reported payments in large bins uninformative for analyses (e.g. \$1-\$1,000, \$1,001-\$10,000, etc.)

[47] Some firms utilized three mutually exclusive categories (e.g., consulting, meals, research), while others utilized non-exclusive labels (e.g., meals; meals, consulting; consulting, teaching and education).

Table A2: Sample Descriptions

		(1)	(2)	(3)	(4)
Panel (a): Unit-of-analysis (*djt*, all drugs)—Claims data					
Cardiologists	unique	19,817	14,449	14,449	13,793
Cardiov. claims	total	1.3e+08	1.1e+08	1.1e+08	1.0e+08
	mean	3,318	3,662	3,662	3,771
	median	2,581	3,002	3,002	3,123
Cardiov. share, focal statins	mean	0.197	0.187	0.187	0.190
	median	0.190	0.187	0.187	0.189
Cardiov. share, Crestor	mean	0.028	0.026	0.031	0.031
	median	0.020	0.021	0.025	0.025
Cardiov. share, Lipitor-2011	mean	0.044	0.041	0.044	0.045
	median	0.035	0.036	0.038	0.038
Cardiov. share, atorvastatin-2012	mean	0.055	0.050	0.052	0.053
	median	0.047	0.046	0.047	0.048
Cardiov. share, other generic sum	mean	0.115	0.110	0.111	0.111
	median	0.108	0.108	0.108	0.108
N *djbt* obs.	unique	217,987	158,939	123,809	121,121
Panel (b): Cardiologist-molecule (*dj*, Crestor or Lipitor)—Payment data					
Cardiologists	unique	19,817	14,449	13,933	13,775
AstraZeneca – Crestor	Any type	0.518	0.593	0.721	0.722
	Total $ amount	485	435	435	437
	Any meal	0.511	0.585	0.712	0.713
	Meal $ amount	84.6	81.3	81.3	81.5
Pfizer – Lipitor	Any type	0.293	0.326	0.350	0.352
	Total $ amount	295	280	280	281
	Any meal	0.274	0.305	0.328	0.329
	Meal $ amount	53.8	51.2	51.2	51.3
N *dj* obs.	unique	39,634	28,898	25,323	25,156
Panel (c): Cardiologist (*d*, Crestor and/or Lipitor)—Payment data					
Cardiologists (N *d* obs.)	unique	19,817	14,449	13,933	13,775
All Types	either firm	0.580	0.660	0.685	0.689
	both firms	0.231	0.258	0.268	0.271
	$ sum	347	362	375	379
Meals	either firm	0.566	0.646	0.669	0.674
	both firms	0.218	0.244	0.253	0.256
	$ sum	57.9	63.2	65.5	66.1

Notes: Reports select summary statistics for prescribing- and payment-related outcomes at three levels of observations (Panels (a–c)) and across four samples. Panel (a) describes prescribing for the full set of doctors and drugs as the data is used for the price and nest regressions. Panels (b) and (c) describe payments from the two branded manufacturers (AstraZeneca–Crestor, Pfizer–Lipitor) at either the doctor-drug level (Panel (b)) or aggregated to the doctor level (Panel (c)). Column (1) obs. based on all cardiologists in Physician Compare; (2) restrictions include those in (1) and then only those for cardiologists with $\geq$ 500 cardiovascular claims in both 2011 and 2012; (3) restrictions include those in (2) and only non-zero q_{djt}; (4) restrictions include those in (3) and then only observations for cardiologists we estimate ψ_{dj} for, which are used in the MTE estimation.

B.3 Comparing the Dataset to Post-Sunshine Act Data

As outlined in the main text, the Kyruus-developed physician-industry interaction data we analyze was available due to the fact that Pfizer and AstraZeneca, among other drug firms, released this information prior to the mandatory reporting regulations of the Sunshine Act, which began reporting in late 2013. Because these disclosures prior to the Sunshine Act

occurred on an ad hoc basis without any standardized reporting agency (the interaction files were typically posted on each firm's website), it is important to provide evidence that this pre-Sunshine Act data is relatively accurate, e.g. it is not censored or biased in any way that would alter our conclusions. To investigate this, we explored post-Sunshine Act data made available by ProPublica[48], examining trends and distributions under the working assumption that firm-level annual trends in physician payments should be smooth, and within-year distributions of payments should be relatively stable.

Like our Kyruus-developed data, the ProPublica version of the official Sunshine Act data, (available at https://openpaymentsdata.cms.gov), is matched to National Provider Identifiers. This enables us to hold fixed our set of cardiologists from the main analyses, and compare payments from Pfizer and AstraZeneca in 2011–2012 (from Kyruus) to those in 2014-2015 (from ProPublica).

Figure A1: Kyruus vs. OpenPayment Comparison

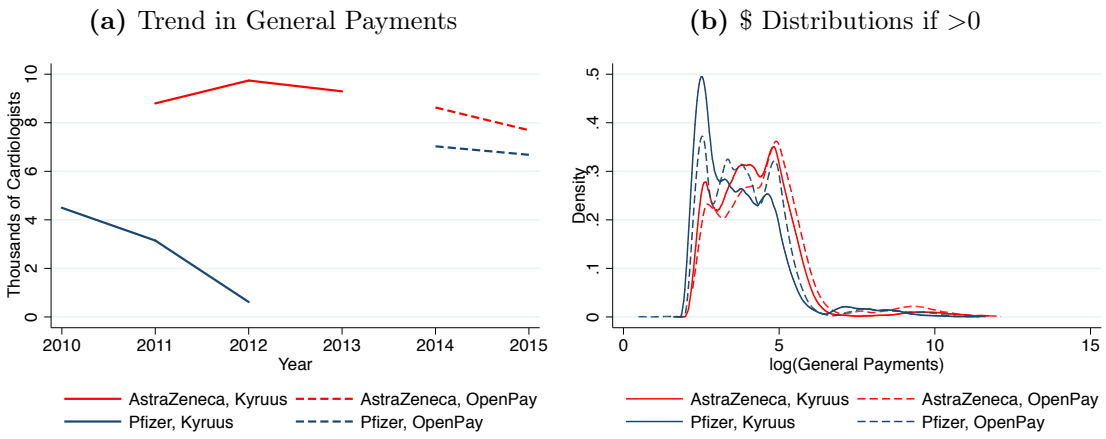

Notes: 2013 is omitted for OpenPayments because OpenPayments reporting only includes the last quarter of that year.

Figure A1 Panel (a) plots the total number of our cardiologists (out of roughly 15,000) that receive any general (non-research) payment from the two firms in each year, based on either data source. In the case of AstraZeneca, the trend is clearly smooth between the two data sources, supporting our assumption that the self-reported data is not notably censored in any way. Although the Pfizer trend line appears to be dramatically different across the two data sources, the spike in 2014 can be explained by the fact that this year marked the approval of Eliquis, a joint venture between Pfizer and Bristol Myers Squibb. Eliquis is an anticoagulant for the treatment and prevention of deep vein thrombosis and pulmonary

[48]https://www.propublica.org/article/about-the-dollars-for-docs-data

embolisms, thus cardiologists are the most relevant specialty, and in the OpenPayments data—where, unlike in the Kyruus data, the specific drug associated with each interaction is reported—Eliquis accounts for roughly 60 percent of the payments to cardiologists and 78 percent of total spending on cardiologists. Figure A1 Panel (b) indicates very little variation in the distribution of payment dollar values across the data years/data sources, further supporting the notion that our data is not censored or biased in any significant way.

C Additional Theory and Connection to Empirics

C.1 Nash Bargaining Solution

We assume that prices of substitute drugs in the market are determined in a simultaneous Nash Equilibrium of Nash Bargaining between suppliers (manufacturers/wholesalers/pharmacies) and buyers (PBMs/insurers). In the model, each price maximizes the Nash Product of the gains from trade for each supplier and buyer pair, taking other prices as given. The first-order condition on each price is:

$$\begin{aligned} p_{jbt}^{pos} &= \arg\max \left(\pi(p_{jbt}^{pos}, p_{jbt}^{oop}, m_{jbdt})\right)^{b_{jbt}} \left(\widetilde{CS}_t(\mathcal{J}_t) - \widetilde{CS}_t(\mathcal{J}_t \setminus jb)\right)^{1-b_{jbt}} \\ &= \left(mc_{jbt} + b_{jbt}\left[\left(1 + \frac{\partial q_{jbt}}{\partial p_{jbt}^{oop}} \frac{p_{jbt}^{oop} - mc_{jbt}}{q_{jbt}}\right) \frac{\widetilde{CS}_t(\mathcal{J}_t) - \widetilde{CS}_t(\mathcal{J}_t \setminus jb)}{q_{jbt}} + p_{jbt}^{pos}(1-\tau_{jt}) - mc_{jbt}\right]\right)/(1-\tau_{jt}) \end{aligned}$$

where $q_{jbt} := \sum_d q_{jbdt}$ denotes the sum over physicians. The term b_{jbt} is a bargaining ability parameter, weighting the extent to which the optimal price depends on supplier profits vs. the expected additional buyer surplus in the case that a contract is agreed to for product jb: $\widetilde{CS}_t(\mathcal{J}_t) - \widetilde{CS}_t(\mathcal{J}_t \setminus jb)$. τ_{jt} reflects the rebate rate to insurers off the posted price p_{jbt}^{pos}.

C.2 Meal Payments: Intuition

Here, we provide a model of the decision by a given drug manufacturer to supply a meal to a given doctor. This model conditions on a global optimization of how to budget meals and the salesforce to execute them across geographic space. As neither our estimation strategy nor our counterfactuals will require solving that global problem, we do not consider it here. Given that global allocation, product j's sales representative should supply a meal to doctor d if the return on investment exceeds whatever hurdle rate R_j the firm applies, which is if and only if:

$$(p_{jr}^{mfr} - mc_j)\left(E[q_{dj}^{m_{dj}=1} - q_{dj}^{m_{dj}=0} | \mathcal{I}_{dj}]\right) > R_j\left(C_{jr}^{m_{dj}=1} - C_{jr}^{m_{dj}=0}\right). \tag{10}$$

Here we assume that the manufacturer price in a region will not change with a meal supplied to one more physician. The key terms are then what the sales representative expects to happen to quantity, given her information set $\mathcal{I}_{dj}$, and the effect of the meal (both direct and indirect) on total costs in the region.

The institutional details in this setting suggest that the cost function $C_{jr}^{m_{dj}=1}$ will have increasing returns to scale in the sense that the average cost of providing a meal will be decreasing in the total meals provided in a region. We would also expect the cost function to depend on other regional characteristics such as the density of candidate physicians in

geographic space. Further, the incremental cost of providing a meal to doctor d is likely to depend on characteristics of that doctor or her employer that affect her willingness to accept a meal.

The expected quantity increase from the meal $E[q_{dj}^{m_{dj}=1} - q_{dj}^{m_{dj}=0} | \mathcal{I}_{dj}]$ will be a function of the expectation of total size of the doctor's patient flow Q_{dt} and the choice probability function as given in (2). In particular, it will be a function of the expectation of the parameter θ_{dj}^m which determines the effect of the meal interaction on the mean utility the doctor assigns to product j.

C.2.1 Meals Equation—Mapping Theory to Empirics

Here, we show how the above theoretical model of meal provision can be simplified to motivate the first stage specification and variables included in our instrumental variables analysis.

We specified that a doctor d would receive a meal from product j whenever

$$(p_{jr}^{mfr} - mc_j)\left(q_{dj}^{m=1} - q_{dj}^{m=0}\right) > C_{dj}^{m=1}(N_{jr_d}, \phi) - C_{dj}^{m=0}(N_{jr_d}, \phi). \tag{11}$$

To deconstruct this expression, we use $\partial q / \partial 1_{\{m>0\}}$ as an approximation to $\left(q_{dj}^{m=1} - q_{dj}^{m=0}\right)$.[49]

We also specify a particular cost function $C_{dj}(N_{jr_d}, \phi) = \phi A_{dj}^{-1/\phi} N_{jr_d}^{1/\phi}$. Here A_{dj} represents an access cost shifter that may be product-doctor specific, N_{jr_d} represents the number of other doctors accessed in the region near d, and this function has increasing returns to scale (decreasing marginal costs of access) iff $\phi > 1$. Here we also use $\partial C / \partial N$ as an approximation to $C_{dj}^{m=1}(N_{jr_d}, \phi) - C_{dj}^{m=0}(N_{jr_d}, \phi)$.

Substituting these values gives

$$(p_{jr}^{mfr} - mc_j) Q_{dj} \frac{\partial s_{dj}}{\partial 1_{\{m_{dj}>0\}}} > A_{dj}^{-\frac{1}{\phi}} N_{jr_d}^{\frac{1-\phi}{\phi}}. \tag{12}$$

Taking logs and rearranging yields a relationship that maps rather cleanly into our linear first stage meals equation:

$$\underbrace{\ln(Q_{dj}) + \frac{1}{\phi}\ln(A_{dj}) - \frac{1-\phi}{\phi}\ln(N_{jr_d})}_{\substack{f(X_{dj};\beta^x) + g(Z_{dj};\beta^z) \\ \text{flexible approx. via Lasso}}} + \underbrace{\ln(p_{jr_d}^{mfr} - mc_j) + \ln\left(\frac{\partial s_{dj}}{\partial 1_{\{m_{dj}>0\}}}\right)}_{\substack{\mu_{dj} \\ \text{residual: correlated with } \theta_{dj}^m + \xi_{dj}}} > 0. \tag{13}$$

[49] For our primary demand specification, this partial derivative is given by: $Q_{dj}\theta_{dj}^m s_{dj}\left(s_{dj} + s_{dj|g}\frac{\lambda}{1-\lambda} - \frac{1}{1-\lambda}\right)$

C.3 Payments in the Presence of Decision Errors

A large literature in economics and health services research has documented that healthcare utilization can be biased upward or downward due to frictions like physician information/skill, physician agency, present bias, symptom salience, and false beliefs. Thus, in order to understand the welfare impact of payments, it is useful to simultaneously characterize the optimality of utilization at baseline, absent payments. Figure A2 below attempts to do this in a very parsimonious way. Consider the mean decision utility term δ_{djt} specified in equation (1). In addition to the impact of meals, this term may include a decision error ε^{de} that shifts perceived quality and, in turn, utilization upward (if $\varepsilon^{de} > 0$) or downward (if $\varepsilon^{de} < 0$).

Figure A2: Role of Payments in the Presence of Decision Errors

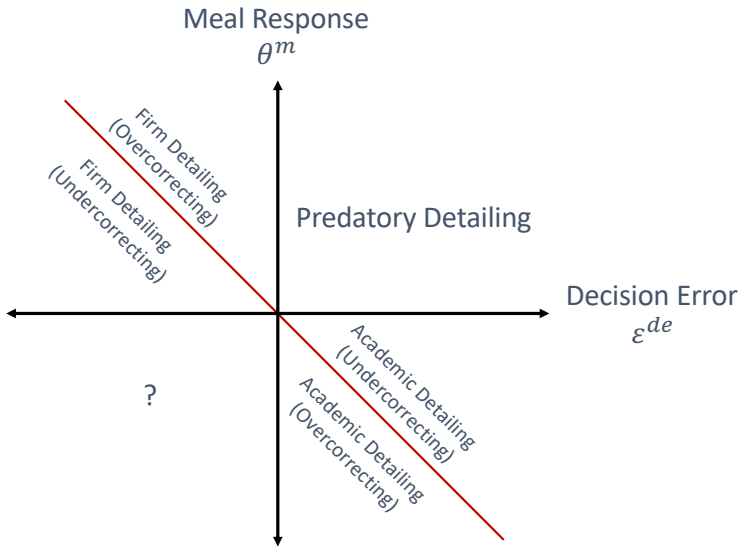

In Figure A2, ε^{de} is on the horizontal axis, while our meal effect θ^m is on the vertical axis. In this framework, several cases emerge.[50] In the upper-right quadrant, both $\varepsilon^{de} > 0$ and $\theta^m > 0$, implying that the firm's detailing efforts exacerbate physicians' existing tendency to overprescribe the drug. For this range of parameters, detailing is predatory and unambiguously inappropriate. In the lower-right quadrant, detailing actually counteracts overprescribing. It would be unexpected for *firms* to voluntarily seek to reduce demand for their own products, but this framework can accommodate a broader range of activities that can be spearheaded by firms, researchers, advocates, or policymakers. These efforts are commonly known as "academic detailing" and may correct utilization downward; see, e.g.,

[50] We are not aware of interventions in the lower-left quadrant (firm or other entities' efforts to reduce utilization for products that are already under-utilized) and will ignore that case.

Midboe et al. (2018) regarding the use of academic detailing to combat the opioid epidemic.

The upper-left quadrant may be most appropriate for understanding firm marketing of statins. As discussed in Baicker et al. (2015), the clinical literature indicates that usage of statins for the indicated patient population is too low, suggesting $\varepsilon^{de} < 0$. However, we estimate that firm efforts push statin use upward significantly: $\theta^m > 0$. Given only price, quantity, and payment data, we cannot directly estimate whether firm efforts result in overcorrecting ($\theta^m > -\varepsilon^{de}$) or undercorrecting ($\theta^m < -\varepsilon^{de}$) among treated physicians and promoted drugs, and this is crucial for policy evaluation. As discussed in Section 5 in the main text, we address this gap by directly incorporating this term ε^{de} into our measure of consumer surplus. Then, in analyzing the welfare impact of a counterfactual ban on meals, we compare a threshold value—the point at which consumer surplus with meals crosses consumer surplus without meals—to an estimate of ε^{de} based on the dollarized health benefit of statins among indicated patients drawn from the clinical literature.

D Parameter Estimation Routine

The following outline details the steps necessary to recover the demand parameters (θ^p, λ, ψ_{dj}, θ^m) and is followed by more in-depth discussions of the Lasso approach we use (Appendix D.1), how the MTEs are estimated (Appendix D.2), the important variables selected by the Lasso algorithm (Appendix D.3), and the role of the perturbation and shrinkage procedures (Appendix D.4).

Jackknife doctors and create bootstrap samples k

1. Replicate the full sample of djt-level observations 500 times, dropping the observations for a randomly selected $\sqrt{N_\text{d}}$ doctors; denote each of these samples k

For each k, perturb quantities

2. Reshape the data to the "use-case" level with a dummy variable $c = 1$ indicating each use (e.g., if $q_{djt} = 50$, this would translate to 50 rows of $c = 1$ use-cases for that djt)
3. Sample with replacement
4. Calculate perturbed quantities $\tilde{q}_{djt} = \sum c_{djt}$

For each k, estimate price, nest and ψ_{dj} parameters

5. Estimate Eq. 4 to recover price ($\theta^{p,k}$), nest (λ^k), and doctor-molecule fixed effects (ψ^k_{dj})
6. Parameter estimates: for θ^p and λ, point estimates and standard errors are given by the mean and standard deviation across the 500 k samples
7. Shrink each ψ^k_{dj} estimate towards the j-specific mean using the standard deviation of ψ^k_{dj} across the 500 k samples as the standard error in the standard empirical Bayes shrinkage formula.

For each k, estimate meal parameters

8. Keep ψ^k_{dj} estimates for Crestor and Lipitor observations
9. Follow the split-sample Lasso approach described below in Section D.1 to select the relevant controls (X) and instruments (Z)
10. Estimate MTEs for meal receipt, $\theta^{m,k}$ as described below in Section D.2 and based on Eqn. 5
11. Parameter estimate: for θ^m (and all corresponding MTE-based estimates) point estimates and standard errors are given by the median and "median deviation" across the 500 k samples as described below in Section D.1

D.1 Split-sample Lasso Approach

Our use of the Lasso draws heavily on Belloni et al. (2017) and Chernozhukov et al. (2018). The outline of our approach, used within each of the k bootstrap samples described above, is as follows:

1. Randomly split the sample into two sub-samples $s = \{A, B\}$
2. Within each sub-sample s, perform a lasso regression of the (dependent variable) ψ on the vector of possible controls (and transformations thereof) X, with the selected variables given by $L(X)^{\psi,s}$
3. Within each sub-sample s, perform a lasso regression of the (endogenous variable) $1_{(m>0)}$ on the vector of possible controls (and transformations thereof) X and possible instruments Z, with the selected variables given by $L(X)^{m,s}$ and $L(Z)^{m,s}$, respectively.

4. Within each sub-sample s, estimate $\theta^{m,k,s}$ (and other MTE parameters) using the variables selected in the opposite sub-sample s' taking the union of $L(X)^{\psi,s'}$ and $L(X)^{m,s'}$ for controls
5. Solve for the k-specific estimate $\theta^{m,k} = (\theta^{m,k,A} + \theta^{m,k,B})/2$
6. Parameter estimate: for θ^m (and all corresponding MTE-based estimates) point estimates and standard errors are given by the median, $\overline{\theta^m} = median(\theta^{m,k})$, and the "median deviation", s.e.$(\overline{\theta^m}) = \sqrt{median(\theta^{m,k} - \overline{\theta^m})}$, respectively.

To minimize functional form assumptions about how the controls enter these functions, squared and log transformations of all X variables are included as possible controls. To allow firm firm-specific responses to the instruments, all Z variables are interacted with molecule-specific dummies.[51]

All Lasso regressions use common machine learning protocols. We use 10-fold cross-validation—split data set into 10 equal parts, and use each in turn as the holdout sample on which the model trained on the other 9 is tested—at 100 potential penalty parameters to select the simplest model (i.e., the largest penalty) that minimizes the mean RMSE in the hold-out samples of the 10-fold cross validation runs. The 100 potential penalty parameters range up to a maximum of $MaxPenaltyGuess = 2 \times \max(\tilde{x}'y)$, where $\tilde{x}$ is the pre-standardized regressor matrix and y is the vector of the outcome variable, from a minimum of $[MaxPenaltyGuess/1000]$; the 100 potential penalties are evenly spaced between the minimum and maximum penalty guess values over a log scale. Our preferred Lasso specification is a "two-step adaptive" model that performs one Lasso, followed by another where only variables selected in the first Lasso are possible controls in the second. Appendix G.6 shows that our results are not sensitive to minor tweaks to this approach.

D.2 MTE Estimation Approach

We estimate MTEs using the mtefe package in Stata 16 (Andresen 2018). Andresen (2018) provides a useful overview of the MTE literature (e.g., Heckman et al. 2006; Heckman and Vytlacil 2007; Brinch et al. 2017) and describes the approach to estimating MTEs that we employ. Briefly put, and borrowing closely from Andresen (2018)'s description, one begins with a generalized Roy selection model, with i indexing individuals, Y denoting potential outcomes D denoting realized treatment, d denoting potential treatments, with W and V

[51] We obtain qualitatively similar results if only using the levels of X variables as controls and/or not allowing instruments to be molecule-specific.

denoting unobservables in the outcome and treatment equations, respectively:

$$\begin{aligned} Y_i^d &= f^d(X_i) + W_i^d \qquad \text{for } d = 0, 1 \\ Y_i &= D_i Y_i^1 + (1 - D_i) Y_i^0 \\ D_i &= \mathbf{1}\{g(X_i,\, Z_i) > V_i\}. \end{aligned} \tag{14}$$

Then we make the two necessary assumptions of conditional independence ($W^d, V \perp Z \,|\, X$: the error terms in the outcome and treatment equations are orthogonal to the instruments conditional on the controls) and separability ($\mathbb{E}[W^d \,|\, V, X] = \mathbb{E}[W^d \,|\, V]$). Per this model and assumptions, MTEs are then defined as:

$$\begin{aligned} MTE(x,\, u) &\equiv \mathbb{E}[Y^1 - Y^0 \,|\, X_i = x, U_i = u] \\ &= \underbrace{x(\beta^1 - \beta^0)}_{\text{heterogeneity in observables (levels)}} + \underbrace{\mathbb{E}[W^1 - W^0 \,|\, U_i = u]}_{\text{heterogeneity in unobservables (slopes)}} \end{aligned} \tag{15}$$

where U, the unobserved resistance to treatment, is given by the quantiles of V.[52]

We encourage the interested reader to see Andresen (2018) for a step-by-step process of the MTE estimation routine via the "separate" approach first outlined by Heckman and Vytlacil (2007). Two specification choices of note: (1) we estimate the propensity scores (meal probability as a function of X and Z) using a LPM since the large number of covariates often led to nonconvergence of probit and logit models; and (2) we use a nonparametric local linear function to estimate the control functions in the model (which are related to $\mathbb{E}[W^1 - W^0 \,|\, U_i = u]$).

As shown by Andresen (2018), posterior estimates of doctor-specific treatment effects can be calculated using the following formula:

$$\begin{aligned} \mathbb{E}[Y_i^1 - Y_i^0 \,|\, X_i = x, D_i = d, P_i = p] = \; & x(\beta^1 - \beta^0) \\ & + d\mathbb{E}[W^1 - W^0 \,|\, U_i \leq p] \\ & + (1 - d)\mathbb{E}[W^1 - W^0 \,|\, U_i > p], \end{aligned} \tag{16}$$

where P_i is the doctor's propensity score.

D.3 Important Variables & the Importance of Many Variables

Table A3 reports some the top variables selected by the Lasso algorithm based on the number of subsamples within which the variable is selected in either the outcome or treatment

[52] This smoothing creates the unit interval that is the x-axis for all MTE curves.

Lasso. With 500 bootstrap samples each with two split-samples, the total possible number of selections is 1,000. The table reports only the top 15 controls (X) and top 10 instruments (Z).

Table A3: Frequently Lasso-selected Variables

	Num. subsamples selected in
Controls (X)	
log(Own cardiov. claims)	998
log(HSA Cardiol. annual claim avg. cardiov.)	981
log(HSA Uninsured share)	976
log(Hosp. admissions)	968
HRR Share faculty	964
Zipcode-drug local market ad spend^2	953
HRR Num. AAMCs^2	922
log(HSA Medicare Advantage eligbl.)	917
log(HSA Doc. annual claim avg. total)	902
log(HSA Doc. annual claim avg. cardiov.)	895
log(HRR Num. doctors)	863
Hosp. Doc. annual claim avg. cardiov.	859
HSA Cardiol. annual claim avg. cardiov.	840
log(HRR Num. cardiol.)	804
log(HSA Medicaid Share)	795
AMC AMSA Instruments (Z)	
Lipitor, HSA AMSA	835
Lipitor, HRR AMSA	822
Crestor, HSA AMSA × faculty-wgt.	674
Lipitor, HSA AMSA × faculty-wgt.	656
Crestor, HRR AMSA × faculty-wgt.	552
Crestor, HSA × faculty-wgt. × drive time	373
Crestor, HRR × drive time	360
Lipitor, HRR × faculty-wgt. × drive time	330
Crestor, HSA AMSA	318
Lipitor, HSA × faculty-wgt. × drive time	310

Notes: Reports the number of subsamples the covariate is selected in either the outcome or treatment Lasso regression; the maximum is 1,000.

Obtaining precise MTE estimates requires substantial overlapping (and preferably full) support of treatment propensities for both the treated and untreated sub-samples. It also requires significant variation in the instrumental variables at each neighborhood on this support in order to estimate the semi-parametric IV regressions. Comparing the common support of treatment propensities under specifications that only use a small subset of covariates in Figure A3 clearly illustrates the value of the Lasso-based approach in that it allows us to generate greater overlapping support and more precision in the MTE estimates.

D.4 Perturbation and Shrinkage

We are concerned that the doctor-molecule mean utility parameters (the ψ_{dj} fixed effects) might be influenced by noise (since we only observe two years of utilization), especially for

Figure A3: Common Supports of Meal Propensity Regression using Different Control Sets

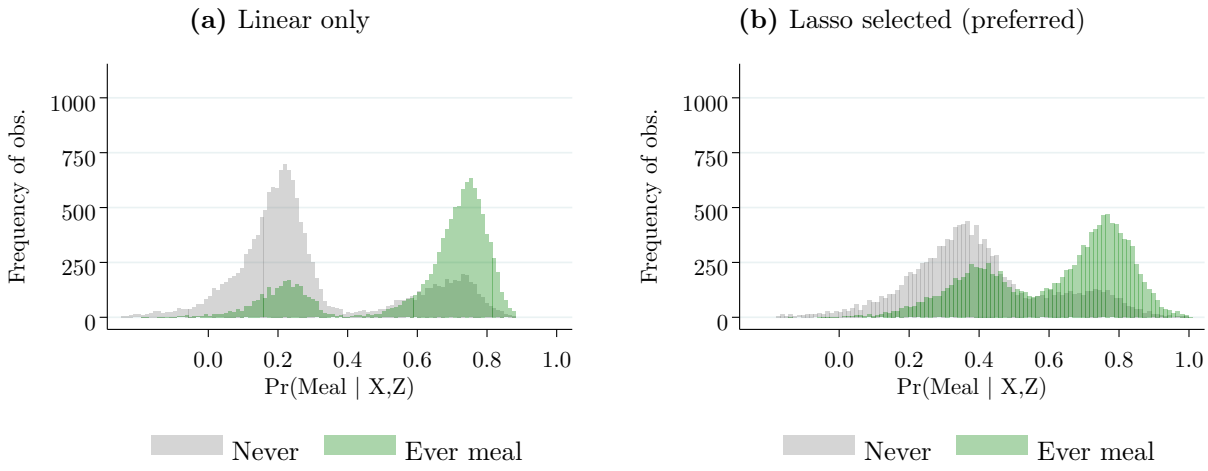

Notes: Meal propensity scores based on the instrumental variables, molecule dummies, and (a) linear versions of the physician-, hospital-, and regional-level covariates, and (b) the Lasso-selected subset of the linear, squared, and log transformations of all covariates (our preferred specification).

low-quantity prescribers. This motivates a "quantity perturbation" procedure. We then use the standard empirical Bayes shrinkage procedure (cf. Chandra et al. 2016) to account for potential estimation error driven by sampling variation.

We used a delete-120 jackknife bootstrap, blocked at the cardiologist level to allow for arbitrary correlations within cardiologist, where we remove 120 physicians (which is the square root of the number of physicians in our sample) from each bootstrap sample. We further resample at the patient level to account for sampling error in market shares. For each subsample, we also follow the sample splitting procedure outlined in Chernozhukov et al. (2018) to prevent contamination of our parameter estimates by overfitting in the machine learning model.

Looking at the raw estimates of ψ_{dj} implied by the price and nest regression with naive physician-level clustered standard errors (Figure A4 Panel a), we observe a counter-intuitive pattern of more precision (larger t-statistics) on ψ_{dj} estimates for low-quantity prescribers. When using the corresponding standard errors from these raw estimates in the shrinkage procedure (Figure A4 Panel b), virtually no adjustment to the estimates occurs. However, when we employ the quantity perturbation we see a more intuitive pattern of precision that is increasing in usage (Panel a), and this introduced variation results in significantly more shrinkage of the estimates. Still, not that much shrinkage occurs; in the case of the perturbed

data, about 94 percent of the observations are shrunk by less than 10 percent of their raw values.

Figure A4: Role of Quantity Perturbation and Fixed-effect Shrinkage

(a) t-statistics for ψ_{dj} per true quantities

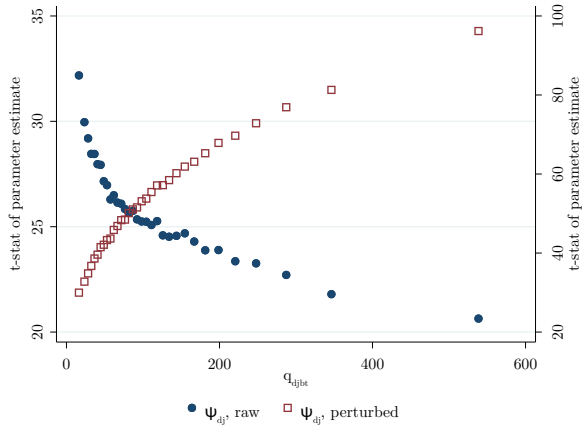

(b) Relative empirical Bayes shrinkage of ψ_{dj}

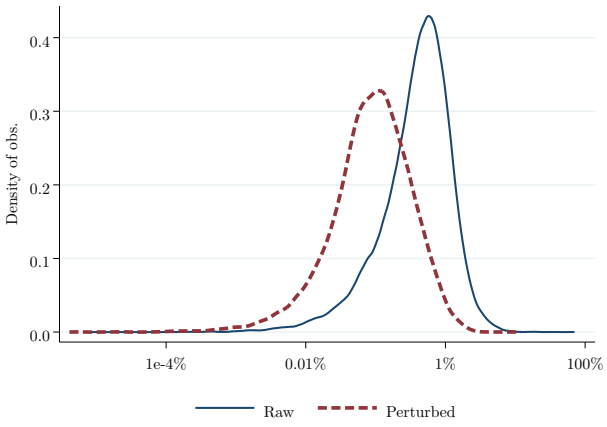

Notes: (a) binned scatterplot of $\psi_{dj}/s.e.(\psi_{dj})$ for the raw data alongside the synthetically perturbed data, where resampling at the claim-level is performed; 33 bins are reported so each marker represents ∼3 percent of the data. (b) Letting $\widehat{\psi_{dj}}$ denote the empirical Bayes estimate of ψ_{dj}, this plots the "relative shrinkage" defined as: $(\psi_{dj} - \widehat{\psi_{dj}})/\psi_{dj}$; note the logarithmic x-axis scale. All data for both panels is based only on the strong statins (Crestor and Lipitor/Atorvastatin).

E Role of Rebates

The negotiation modeled in Section 5.3 is described as taking place between an abstract "supplier" and "buyer." The pharmaceutical supply chain is complex, in that there are both supply (wholesalers, distributors, pharmacies) and demand (PBMs) intermediaries with market power, and multiple bilateral negotiations take place between these parties (The Health Strategies Consultancy LLC 2005). Ultimately, we only observe the point-of-sale price paid by buyers when prescriptions are filled—we do not observe confidential rebates remitted back to insurers/PBMs, and we do not observe the unit price paid directly to manufacturers. In practice, we account for these issues using average data on rebates and intermediary profits.

This assumption comes to bear in two parts of our analysis. First, the prices p^{pos}_{jrt} that suppliers receive, and that insurers pay after cost-sharing is applied, are net of rebates τ. This is an approximation, as we are collapsing a set of bilateral negotiations between upstream and downstream firms into a single negotiation over a unit price, and the "producer surplus" is split between manufacturers, wholesalers, distributors, and pharmacies. Second, panel (c) of Figure 4 plots the distribution of expected changes in firm revenue from targeting meal payments. Naturally, we expect drug manufacturers to determine meal targeting as a function of their own revenue only. Thus, in this analysis, we allow for unobserved rebate τ *and* "other suppliers' markup" w, so that manufacturer revenue becomes $R(p^{pos}_{jrt}) = \sum_{d \in r} q_{djt} p^{pos}_{jrt}(1 - w - \tau)$.

In the supply side estimation, welfare simulations, and simulations of manufacturer revenue, we rely on researchers' estimates of τ and w, and we test the sensitivity of our results to our decisions on how to use these estimates.

We are aware of several sources of information on $\tau + w$: Yu et al. (2018) use 2016 list price and net price estimates from IQVIA. IQVIA's estimates are themselves based on manufacturers' filings with the Securities and Exchange Commission (SEC), publicly reported net sales, and information provided by these companies directly in support of IQVIA's analysis, for a large sample of pharmaceutical companies. Kakani et al. (2020) use similar data from SSR Health, LLC going back to 2012. Sood et al. (2017) report data collected directly from sources such as SEC filings. In each case, the researchers report prices obtained *by manufacturers* after rebates, discounts, concessions, etc. The results are very similar: Yu et al. (2018) reports an overall net price of $p^{mfr} = p^{pos} * (1 - w - \tau) = 0.673 * p^{pos}$, suggesting $\tau + w = 0.327$. Kakani et al. (2020)'s estimates suggest an average $\tau + w = 0.32$ across a wide range of drugs that (unfortunately) explicitly excluded statins. Sood et al. (2017) suggests $\tau + w = 0.42$ across a range of branded and generic drugs.

To obtain the components τ and w, we rely on multiple sources. CMS has reported average manufacturer rebate percentages overall ($\tau = 0.175$) and for cardiovascular drugs specifically ($\tau = 0.263$) going back to 2014.[53] Arcidiacono et al. (2013) assume $\tau = 0.151$ and estimate that (in the antiulcer drug market) rebates increase to 48.3 percent after branded drugs' patents expire. Similarly, Aitken et al. (2018) suggest that Lipitor rebates increased in 2012. One can also infer w from Yu et al. (2018), as they pulled together aggregate data on profits to PBMs, wholesalers, pharmacies, providers, and insurers. We ignore the profits of PBMs and insurers, as those are "buyers" in our calculation. We also ignore provider profits, as those refer to physician-administered drugs such as chemotherapy and are not relevant for statins. That leaves wholesalers and pharmacies, which are estimated to obtain profits of $0.037 * p^{pos}$ and $0.152 * p^{pos}$, respectively. Thus, the work in Yu et al. (2018) suggests that $w = 0.190$.

For our simulations of manufacturer revenue, the above papers suggest $\tau + w = 0.32$ if statin markups and rebates look like those of the average drug in the US. If statins instead follow other cardiovascular drugs in having relatively high rebates, then $\tau + w = 0.263 + 0.190 = 0.453$ would be more appropriate. For our supply side estimation and counterfactuals, the above papers suggest $\tau = 0.32 - 0.037 - 0.152 = 0.131$ as a lower bound based on patterns observed for a wide range of pharmacy-dispensed drugs (Kakani et al. 2020; Yu et al. 2018) and $\tau = 0.263$ based on cardiovascular drugs only. We use $(\tau = 0.263, \tau + w = 0.453)$ in the main text and $(\tau = 0.131, \tau + w = 0.32)$ in Table A4 below. These figures refer to the values used for branded drugs pre-patent expiration. For Lipitor in 2012, we decrease the pass-through to Pfizer in the main text by using $(\tau = 0.483, \tau + w = 0.673)$ (based on Aitken et al. (2018); Arcidiacono et al. (2013)); we stick with the alternative assumptions $(\tau = 0.131, \tau + w = 0.32)$ in the Appendix for robustness. Finally, for generic drugs, we rely on Sood et al. (2017), which is the only source explicitly breaking out generics, and assume $(\tau = 0.24, \tau + w = 0.41)$.

Comparing Table A4 to the results from our preferred specification in the main text (Figure 4, Panel c), this alternative (larger) pass-through assumption very intuitively yields larger revenues. But the differences are not substantial, as we cannot statistically reject differences between the two pass-through assumptions at any of the points of the distribution that we report here for either the never- or ever-treated physicians.

[53] https://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/Information-on-Prescription-Drugs/PartD_Rebates

Table A4: Heterogeneity in Revenue Effects, Alternative Rebate Assumption

	$p10$	$p25$	$p50$	$p75$	$p90$
Never	428.8	1,007.1	2,236.7	4,581.0	8,203.4
	(687.6)	(896.9)	(1,509.7)	(2,651.7)	(4,217.7)
Ever	1,260.0	2,221.4	4,327.7	8,098.0	13,702.9
Meal	(255.1)	(431.7)	(809.3)	(1,484.8)	(2,340.2)

Notes: Plots the distribution of marginal revenues due to meals. Revenues are based on our estimate of the net price p^{mfr} paid to manufacturers, net of rebates and markups charged by supply intermediaries. Here, we assume $p^{mfr} = 0.68 * p^{pos}$. Beneath each plot is the point estimates and standard errors for select percentiles of these distributions by treatment status.

F Dollar Value of Health Gains

In this Appendix, we estimate the dollar value of the health benefits of statins based on evidence in the clinical literature. Throughout, we err on the side of making conservative choices when assumptions are necessary with one crucial exception: we assume that patients *are clinically indicated* for statins. This seems reasonable given the American College of Cardiology assertion that full adherence to clinical guidelines would increase statin use by 24 percent relative to baseline (American College of Cardiology 2017).

We take the perspective of a decision-maker deciding whether to have an indicated patient initiate statin therapy given the expected health benefits and out-of-pocket costs. Unfortunately, many individuals initiating a medication regimen do not adhere to that regimen long enough to experience health benefits. For our analysis, we assume 37 percent adherence at five years, which is the bottom of the range in Deichmann et al. (2006)'s meta-analysis and is close to the adherence level implied by Colantonio et al. (2019) (78 percent adherence year over year for five years) and Colantonio et al. (2017) (40 percent of statin initiates seeing a full 5-year benefit).

The Heart Protection Study Collaborative Group indicates a benefit (of taking a statin, vs. nothing) of about 0.69 life years for Medicare-age enrollees if adherence is perfect over five years; the estimated benefit drops to 0.31 life years if adherence declines to 35 percent by the sixth year (Heart Protection Study Collaborative Group (2009)). Based on this, we make two conservative assumptions. First, we assume that 37 percent of patients initiating therapy under a given regime are perfectly adherent and receive health benefits; all others receive no health benefit. Second, we focus on the benefits of expanding statin use overall, so we do not differentiate generic statins and "strong statins," for which there is clinical evidence that strong statins lead to an additional 0.09 life-year gain among indicated patients (see, e.g., Wagner et al. 2009).

Finally, we use a value of $75,000 per life-year gained, which is at the bottom of the $75,000-$100,000 range in Cutler (2004). We do not inflation-adjust, for the sake of simplicity.

Taken together, the above estimates indicate a dollarized health benefit of $0.69 * 0.37 * 75,000 = \$19,147.50$ is associated with initiating a statin regimen. The appropriate out-of-pocket cost comparison is with the total out-of-pocket cost of a statin regimen over five years, in present discounted value and with 78 percent adherence each year. In contrast, the out-of-pocket cost in our demand analysis is out-of-pocket price for a single month. Accordingly, to obtain the "flow" value of taking statins, we divide the dollar value above by the multiplier on monthly out-of-pocket costs that is necessary to cover five years of prescription statins.

Using a 3 percent interest rate, this multiplier is:

$$f_{5yr}^{oop} = 12 \sum_{n=1}^{5} 1 * \left(\frac{0.78}{1.03} \right)^{n-1} = 37.13$$

Thus, the "flow" dollarized health benefit of statins is $\$19,147.50/37.13 = \$516$.

G Additional Tables and Figures

G.1 Summary Statistics

Tables A5 and A6 report the summary statistics and univariate regression coefficients for the 75 variables that form the basis of our potential control set in the Lasso meal regressions. Both tables are based only on doctor-molecule observations for 2011 and either Lipitor (Pfizer) or Crestor (AstraZeneca). Each concentric group excludes observations from within the same preceding group.

Table A5: Summary of Potential Controls—Doctor, Hospital, and Zipcode

	mean	s.d.	β meal	β cardiov. share
Doctor level:				
log(Claims, cardiov.)	7.81	0.858	0.290	–0.191
ERx = 1	0.709	0.454	0.121	–0.049
Num. hospitals	3.64	1.42	0.078	–0.118
Med. school grad. year	1,985	9.86	0.044	–0.062
Num. orgs.	1.54	0.838	0.041	–0.016
EHR = 1	0.632	0.482	0.039	–0.020
Num. zipcodes	1.69	1.74	0.003	–0.017
PQRS = 1	0.534	0.499	0.003	–0.009
Drive time to nearest AAMC	877	1,900	–0.017	–0.053
Num. specialties	1.40	0.550	–0.031	–0.0042
Female = 1	0.088	0.284	–0.096	0.016
AAMC Faculty = 1	0.096	0.295	–0.226	0.120
AMSA	2.27	7.46	–0.234	0.118
Hospital level:				
Cardiol. annual avg. cardiov. claims	2,949	2,109	0.219	–0.116
Cardiol. annual avg. all claims	3,745	2,819	0.211	–0.108
Cardiol. annual sum all claims	53,923	53,353	0.134	0.106
Doc. annual avg. all claims	1,592	940	0.133	–0.157
Cardiol. annual sum cardiov. claims	42,385	40,811	0.132	0.103
Doc. annual avg. cardiov. claims	719	475	0.127	–0.130
Doc. annual sum all claims	506,158	395,076	0.054	0.049
Doc. annual sum cardiov. claims	229,650	180,194	0.052	0.062
AMSA avg.	23.9	7.44	0.003	0.076
Num. beds	357	347	–0.064	0.093
Num. admissions	17,719	17,452	–0.067	0.100
Num. AAMCs	4.41	3.68	–0.070	0.212
Num. cardiologists	18.0	19.7	–0.082	0.204
Num. doctors	390	357	–0.140	0.183
Num. AAMC faculty	61.3	171	–0.252	0.177
Share AAMC faculty	0.081	0.152	–0.280	0.171
AMSA faculty-wgt.	2.11	4.05	–0.285	0.169
Zipcode-drug level:				
Ad spend	33,876	48,266	–0.162	0.261
Ad duration	64,115	23,172	–0.654	0.279
Ad units	1,070	386	–0.656	0.278

Notes: e-RX: electronic prescribing system; PQRS: participates in the Physician Quality Reporting System; EHR: electronic health records. The "β meal" and "β cardiov. share" columns report the coefficient from a regression of either a dummy for meal receipt (β meal) or the standardized cardiovascular share of a drug (β cardiov. share) regressed on the standardized variable.

Table A6: Summary of Potential Controls—HSA, HRR, and State

	mean	s.d.	β meal	β cardiov. share
HSA level:				
Uninsured share	10.7	4.36	0.156	–0.082
Cardiac hospitalizations per 1K	66.4	11.6	0.097	–0.072
Cardiol. annual avg. cardiov. claims	1,468	7,488	0.039	0.0082
Cardiol. annual avg. all claims	1,868	9,458	0.039	0.0091
Medicare Advant. eligibility	118,746	196,160	0.016	0.109
Doc. annual avg. cardiov. claims	432	5,108	0.001	–0.001
Doc. annual avg. all claims	949	12,099	–0.003	–0.001
Cardiol. annual sum all claims	120,896	215,909	–0.007	0.084
Cardiol. annual sum cardiov. claims	93,604	161,980	–0.008	0.079
Medicaid share	22.0	8.50	–0.014	0.055
Doc. annual sum all claims	132,2391	1,994,115	–0.057	0.074
Num. cardiologists	49.4	85.8	–0.059	0.121
Medicare Advant. penetration	23.3	13.7	–0.059	0.022
Doc. annual sum cardiov. claims	579,322	879,976	–0.063	0.076
Num. AAMCs	8.09	11.4	–0.073	0.096
Num. doctors	1,599	2,384	–0.087	0.115
Num. AAMC faculty	282	626	–0.092	0.134
Teaching hosp. admissions share	0.122	0.214	–0.097	0.093
Teaching hosp. bed share	0.112	0.192	–0.101	0.085
Share AAMC faculty	0.036	0.024	–0.120	0.162
HRR level:				
Teaching hosp. bed share	0.145	0.151	–0.051	0.073
Num. AAMC faculty	421	850	–0.065	0.030
Num. doctors	4,592	4,589	–0.078	0.066
Uninsured share	10.7	4.36	0.143	–0.073
Cardiac hospitalizations per 1K	66.5	11.6	0.094	–0.067
Medicare Advant. eligibility	117,983	197,103	0.025	0.100
Cardiol. annual avg. cardiov. claims	1,024	4,753	0.018	0.004
Cardiol. annual avg. all claims	1,332	6,104	0.018	0.005
Doc. annual avg. cardiov. claims	326	2,884	0.014	–0.007
Doc. annual avg. all claims	886	7,926	0.008	–0.006
Cardiol. annual sum cardiov. claims	64,548	114,387	0.001	0.060
Cardiol. annual sum all claims	85,736	158,674	0.001	0.066
Medicaid share	22.0	8.49	–0.011	0.042
Doc. annual sum all claims	1,098,171	1,703,701	–0.042	0.059
Doc. annual sum cardiov. claims	422,471	668,739	–0.044	0.062
Num. cardiologists	48.2	83.8	–0.048	0.104
Medicare Advant. penetration	23.3	13.7	–0.053	0.023
Num. AAMCs	7.98	11.3	–0.066	0.082
Num. doctors	1,566	2,349	–0.072	0.097
Num. AAMC faculty	272	602	–0.074	0.117
Teaching hosp. admissions share	0.119	0.212	–0.080	0.082
Teaching hosp. bed share	0.110	0.190	–0.084	0.076
Share AAMC faculty	0.036	0.023	–0.101	0.133
State level:				
Plan enrollment	992,164	689,303	0.070	0.062
Plan enroll., low income subs.	410,867	301,107	0.069	0.078

*Notes*The "β meal" and "β cardiov. share" columns report the coefficient from a regression of either a dummy for meal receipt (β meal) or the standardized cardiovascular share of a drug (β cardiov. share) regressed on the standardized variable.

G.2 HSA-level First Stage for Meals

Table A7 replicates Figure 2 Panel (d) in the main text, here showing the HSA-level results of the policy spillover first stage regressions. See the table notes for details.

Table A7: HSA-level I.V. First Stage

	(1)	(2)	(3)	(4)
AMSA avg.	–0.0116	–0.0115	–0.0114	–0.0112
	(0.001)	(0.0012)	(0.0011)	(0.0013)
AMSA avg. × time to nearest AAMC		–0.0004		–0.0004
		(0.0007)		(0.0007)
AMSA avg. × faculty-wgt.			–0.0012	–0.0012
			(0.002)	(0.0021)

Notes: reports the OLS regressions of meal indicator on the vector of Lasso-selected controls (of which own- and hospital-level AMSA scores are a potential control) and manually selected HSA-level instruments. All instruments are standardized; the interactions with the driving time to the nearest AAMC are the products of the standardized variables. Point estimates and standard errors are based on the perturbation-bootstrap approach described in the text.

G.3 Placebo Tests of AMSA Instruments

While a direct test of the exclusion restriction is not possible, empirical researchers often employ "placebo" tests where, for reasons unrelated to the focal identification strategy, certain subsamples of individuals are forced into treatment or non-treatment and are therefore immune to the instrumental variables (Bound and Jaeger 2000; Altonji et al. 2005; Angrist et al. 2010). If the researcher can show that, for such subsamples, the "first stage" relationship (effect of the I.V.s on treatment propensity) and the "reduced form" relationship (direct effect of the I.V.s on the dependent variable) no longer exist, this is encouraging evidence that the reduced form relationship in the full sample is not driven by unobservables.

While we do not have a perfect subsample for whom all meal payments are shut down for exogenous reasons, we perform this exercise for four samples that we expect to be particularly insensitive to the AMC policy spillovers we use as instruments. The four samples are defined as follows:

1. Restricted states: cardiologists in Vermont, Minnesota, and Massachusetts, which either had a complete (VT) or partial ban on certain forms of gifts from pharmaceutical firms to physicians (MN, MA).
2. Hi-AMSA Faculty: faculty cardiologists in the top 25 percent of AMSA scores
3. Hi-AMSA Hospital: cardiologists at hospitals in the top 5 percent of faculty-weighted hospital AMSA scores
4. $P(1_{m>0} \mid X) < 0.1$: cardiologists with meal propensity scores (based only on the X controls) below 0.1.

For subsamples (1–3), we expect either the state-based bans or direct institutional policies to have a dominating effect on firms' (representatives') decisions to pursue relationships with the doctors for whom these policies apply. For subsample (4), we leverage our large set of covariates and the support of meal propensities it creates to pursue a generalized version of these placebo tests, looking only at observations who (per their observables) have a very low probability of meal receipt. We examine both cardiovascular shares (s_{djt}) and the doctor-molecule mean utilities defined in Section 4 (ψ_{dj}) as dependent variables.

Since we have multiple possible instruments, the reporting and interpretation of coefficients in this sort of test is difficult. For simplicity, Table A8 reports the results of the aforementioned placebo tests without any Lasso selection involved, using just OLS and only the linear version of the potential control variables and instruments.[54] Instead of reporting the coefficient on every instrument, for each regression we report the average change in predicted meal probability that is associated with a 1 s.d. change in all AMC AMSA scores.

[54] Using this model yields 2SLS estimates very similar to the post-Lasso 2SLS model with the full set of transformed control variables included.

To obtain this, we calculate the magnitudes of the instruments in this counterfactual world where all AMCs have a 1 s.d. larger (more strict) AMSA score, predict their new meal propensity, then take the difference between this new value and the "actual" propensity (per observed AMSA scores), and then take the average of these differences across all observations. The p-values reported in the table are based on the joint test of significance of the instruments in each regression.

Across the four placebo tests we conduct (Table A8, columns 4–12), we see that both the first stage and reduced form effects are significantly reduced compared to what we observe in the full sample (columns 1–3). In all cases, we either cannot reject the null of the instruments having no effect on meal receipt or the dependent variables, or we estimate a precise but very small effect (often an order of magnitude smaller than the full-sample estimates). These results—that, among cardiologists likely unaffected by AMC policy spillovers, these spillovers are not predictive of either meal payments or our dependent variables—provides support for our exclusion restriction.

Table A8: First Stage and Reduced Form Estimates—Full and Placebo Samples

	Full sample			Restricted states			Hi-AMSA Faculty		
	FS	RF: s_{djt}	RF: ψ_{dj}	FS	RF: s_{djt}	RF: ψ_{dj}	FS	RF: s_{djt}	RF: ψ_{dj}
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\frac{\partial P(1_{m>0}) \text{ in p.p.}}{\partial \text{AMSA}}$	–5.93	–0.138	–2.528	–0.256	0.022	0.609	–0.003	–0.005	–0.001
p-value	[<0.001]	[<0.001]	[<0.001]	[0.01]	[0.384]	[0.654]	[0.007]	[0.676]	[0.168]
N obs.		23,392			863			482	
mean $1_{m>0}$ in p.p.		51.8			19.5			25.5	

Table A9 (cont'd)

	Hi-AMSA hospital			$P(1_{m>0} \mid X) < 0.1$		
	FS	RF: s_{djt}	RF: ψ_{dj}	FS	RF: s_{djt}	RF: ψ_{dj}
	(7)	(8)	(9)	(10)	(11)	12)
$\frac{\partial P(1_{m>0}) \text{ in p.p.}}{\partial \text{AMSA}}$	0.040	–0.013	–0.177	–0.031	–0.003	0.061
p-value	[0.308]	[0.025]	[0.612]	[0.168]	[0.125]	[0.329]
N obs.		1,107			514	
mean $1_{m>0}$ in p.p.		28.1			10.0	

Notes: Reports the first stage (FS) and reduced form (RF) OLS regressions using the all instruments and the only the levels of the control variables. For the first stage regressions $\frac{\partial P(1_{m>0}) \text{ in p.p.}}{\partial \text{AMSA}}$ is the change in meal probability (in percentage points for ease of viewing; $\in [0, 100]$) implied by the coefficients on the set of instruments if all AMCs increased their CoI policies by one standard deviation per the AMSA scores. The mean meal probability (also in percentage points) for each sample is reported in the bottom row. In RF: s_{djt}, the dependent variable in the reduced form regressions is the 2011 focal drugs' share of cardiovascular claims, and in RF: ψ_{dj}, the dependent variable in the reduced form regressions is the estimated doctor-molecule intercept from the price and nest regression (ψ_{dj}). As with the first stage, the reduced form effects are multiplied times 100 for ease of viewing. p-values are based on a test of joint significance of the set of instruments. The effect estimates and p-values are based on the same 500 bootstrap samples used in the demand estimation algorithm.

G.4 Alternative Specifications for Price and Nest Regression

Table A9 shows the demand parameter estimates for several different specifications to help to illustrate how our instrumental variables move coefficient estimates and the effects of different nesting structure assumptions. Column (1) replicates our preferred specification with a statin nest, and uses instruments for both the price parameter and the nest parameter. Columns (2–3) instrument only the nest or price parameter, respectively. While they yield relatively similar price elasticities, we estimate noticeably different nest and price parameters that imply significantly different substitution patterns.

Not performing our quantity-perturbation procedure, dropping AMC faculty from these regressions, and not including a statin nest all yield estimates similar to our preferred specification (columns 4–6). A two-level nesting structure with a statin nest and another nest just for strong statins (Lipitor, Crestor, and generic atorvastatin; column 7) yields results very similar to our preferred specification.

Table A9: Alternative Demand Specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
θ^p	–0.00761	–0.00716	–0.00026	–0.00754	–0.00764	–0.01326	–0.00752
	(0.00016)	(0.00013)	(0.00002)	(0.00003)	(0.00017)	(0.0001)	(0.00017)
λ_{statin}	0.421	0.276	0.971	0.422	0.425		0.438
	(0.011)	(0.012)	(0.001)	(0.002)	(0.011)		(0.011)
$\lambda_{\text{strong statin}}$							–0.024
							(0.008)
mean(η^p)	–0.21	–0.16	–0.13	–0.21	–0.21	–0.23	–0.21
s.d.(η^p)	0.23	0.17	0.15	0.23	0.23	0.24	0.23
F–stat.	680.1	1065.3	15038.4	793.2	561.7	45276.0	288.9
mean($\psi_{dj}/\lvert\theta^p\rvert$) $_{\text{str. statins}}$	–294.4	–348.4	–6005.6	–296.7	–294.7	–209.9	–295.5
mean($\psi_{dj}/\lvert\theta^p\rvert$) $_{\text{other gen.}}$	–313.3	–365.7	–6126.5	–315.6	–313.5	–227.3	–313.3
s.d.($\psi_{dj}/\lvert\theta^p\rvert$)	74.5	91.3	1499.1	75.6	73.8	63.3	74.2
$R^2(\delta_{djt} : \psi_j)$	0.282	0.298	0.017	0.287	0.285	0.287	0.284
$R^2(\delta_{djt} : \psi_j + \theta^p p)$	0.417	0.401	0.004	0.425	0.422	0.418	0.418
$R^2(\delta_{djt} : \psi_{dj} + \theta^p p)$	0.802	0.827	0.138	0.787	0.8	0.872	0.797
Specification							
Instrument θ^p	Y		Y	Y	Y	Y	Y
Instrument λ	Y	Y		Y	Y	Y	Y
Perturb q	Y	Y	Y		Y	Y	Y
Drop AMC faculty					Y		

Notes: Replicates the price and nest regression using the preferred specification (Col. 1) and six alternate specifications. Parameter estimates based on Eq. 4. Point estimates are based on the average of the 500 bootstrap samples (also perturbed per specification); standard errors for the main parameters (θ^p and λ), in parentheses, are based on the standard deviation of these 500 point estimates. $R^2(\delta_{djt} : \cdot)$ reports the R^2 from a regression of $\delta_{djt} \equiv \ln(s_{djt}/s_{d0t})$ on some combination of the molecule- (ψ_j) or molecule-doctor-level (ψ_{dj}) fixed effects, and possibly the price effect ($\theta^p p$). The average number of observations in each bootstrap sample is 116,559.

G.5 Exploration of Treatment Effect of Payments

In the following analyses, we explore alternative indicator variables for payments, an alternative approach to capturing business stealing, and the extent of intensive margin (dollar value) meal effects.

Figure A5 plots MTE estimates of θ^m and, in the table below, the ATE and LATE implied by the MTEs, using: our preferred indicator for meal-based relationships (Col. 1, "Ever meal"); an indicator for physicians receiving meals in 2011 (Col. 2); and an indicator for receiving any kind of payment (e.g., meals, consulting, speaking, travel, or research) from the firm (Col. 3). Results are very similar between our preferred specification and the 2011-only specification; results are somewhat smaller and noisier for "all payment types," but are not statistically significantly different from our main specification.

Figure A5: Alternative Payment Indicators

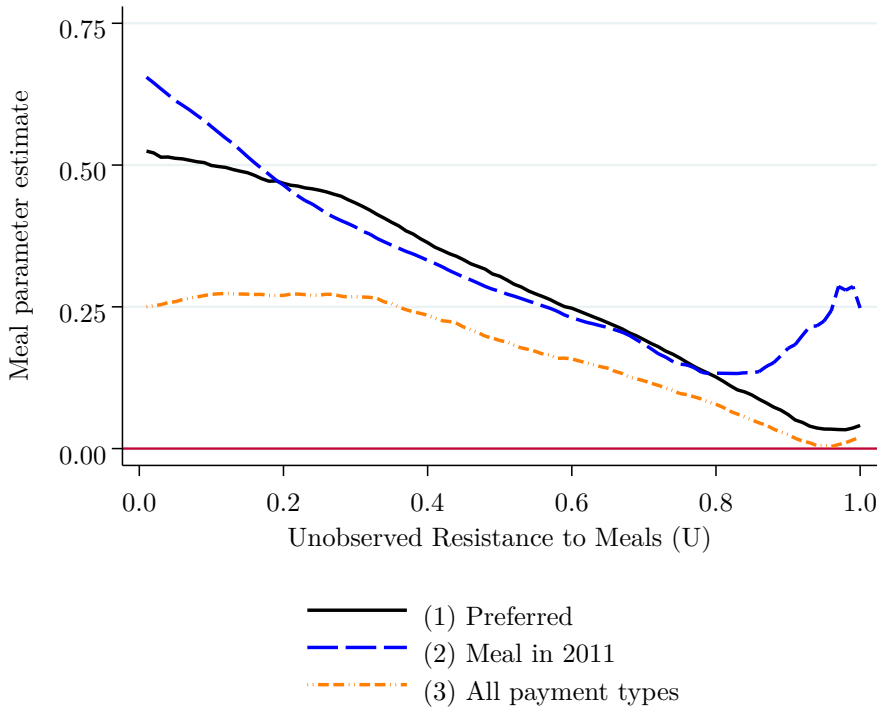

	(1)	(2)	(3)
ATE	0.299	0.329	0.179
	(0.0858)	(0.0819)	(0.1236)
LATE	0.343	0.383	0.242
	(0.0824)	(0.0694)	(0.1246)
F-stat	138.8	184.9	140.9

Notes: The columns in the table correspond to the MTE curves indicated by the legend in the figure; see the accompanying text for details of the three specifications (1–3).

To investigate business stealing and intensive margin effects, we resort to a traditional 2SLS model, since both of these questions involve multiple endogenous variables and estimating MTEs with multiple endogenous variables in a single equation is a task beyond the scope of this paper. As our hypothesis tests of interest are tests of null effects, the 2SLS estimator should still be informative. Columns 1 and 2 of Table A10 report the 2SLS results where we include the same meal indicator from our preferred models, but we also include the average dollar amount of the meal-based payments (per year) as a second endogenous variable, again instrumented by the Lasso-selected CoI policy instruments. Since this is effectively an interaction term (it equals zero for all non-paid physicians), we demean the dollar amounts using either sample-wide (Col. 1) or firm-specific (Col. 2) average dollar amounts. In both cases, we estimate precise zero effects of meal dollar value on prescribing, conditional on the dummy variable for meal receipt.

Lastly, we explore the role of business stealing between the two branded manufacturers in columns 3 and 4 of Table A10. The nested logit model in the main text already incorporates a role for business stealing in the net effects of meal payments; the gains in market share to a paying firm come at the expense of the other products (and the outside good). But the nature of this substitution is strictly defined by the logit functional form. A more flexible specification would allow the receipt of a meal from one branded firm to have a direct effect on the market share of the other branded firm's product (e.g., as in Sinkinson and Starc 2019).

We explore this by adding an additional endogenous variable in the model shown in column 3 that indicates whether the physician was paid by the focal product's rival firm. The model remains identified because we have multiple instruments and we allow each to be interacted with a firm-specific dummy. Theoretically, this specification would be justified by, for instance, differential costs of detailing or strategic responses to AMC effects by the two firms in question. When including this additional term, we estimate a larger own-firm effect of meals (0.66 versus 0.55). The effect of rival firm meals on own-product prescribing is estimated to be negative (–0.21), which would indicate more business stealing than suggested by the nested logit functional form. However, neither of these changes is statistically significant, relative to the baseline specification.

Table A10: Extensive/Intensive Margins and Business Stealing Tests

	(1)	(2)	(3)
Meal, own firm	0.540	0.556	0.663
	(0.1355)	(0.1389)	(0.0911)
Meal $-value,	0.0002	–0.001	
own firm	(0.0020)	(0.0014)	
Meal, other firm			–0.214
			(0.1343)
$-value demean	Full-sample	Firm-specific	

Notes: Replicates the 2SLS-version of the results reported in Figure 3 with additional endogenous variables.

G.6 Alternative MTE Results and Specifications

Figure A6 recreates the MTE curves and displays the ATE / LATE estimates corresponding to our preferred specification shown in the main text (1), as well as five alternative specifications (2–6). Specification (2) includes the hospital-level AMSA controls (Xs) from all of each cardiologists' secondary affiliations (only the AMSA scores of their primary affiliation hospital are used in the preferred model). Specification (3) uses a "one-step" Lasso regression in the variable selection routines (the preferred model uses a "two-step" Lasso). Specification (4) does not perform the empirical Bayes shrinkage (the preferred model does). Specification (5) excludes all AMC faculty from the entire estimation routing (the preferred model includes them). Specification (6) does not perturb the annual claim quantities at the use-case level (the preferred model does).

Figure A6: Alternative MTE Specifications

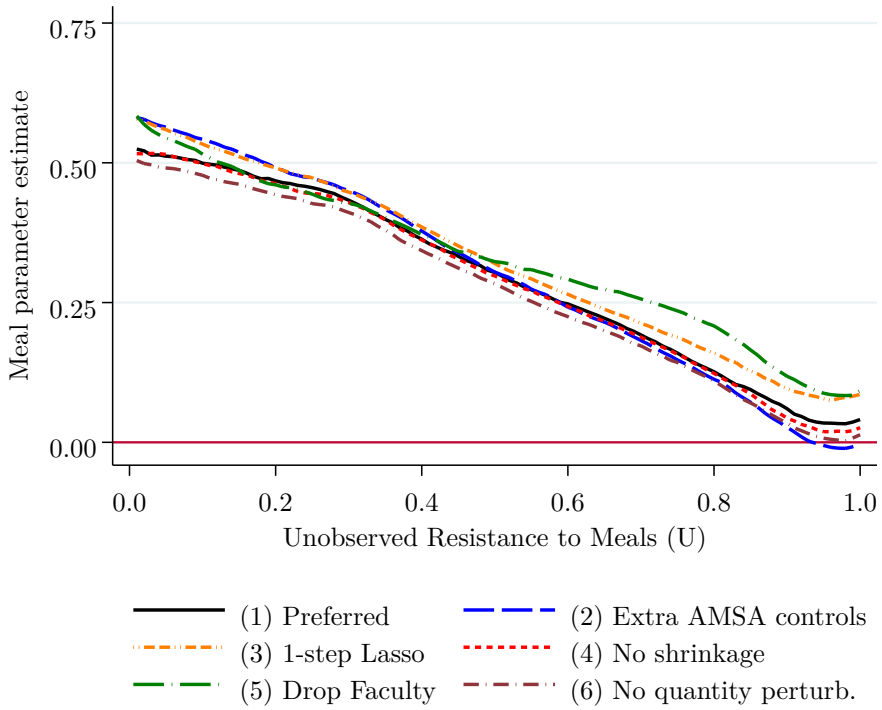

	(1)	(2)	(3)	(4)	(5)	(6)
ATE	0.299	0.305	0.330	0.299	0.333	0.279
	(0.0858)	(0.0873)	(0.0664)	(0.0862)	(0.1001)	(0.0824)
LATE	0.343	0.365	0.376	0.350	0.355	0.330
	(0.0824)	(0.0817)	(0.0586)	(0.0819)	(0.0948)	(0.0779)
F-stat	138.8	144.6	98.3	139.6	154.3	141.5

Notes: The columns in the table correspond to the MTE curves indicated by the legend in the figure; see the accompanying text for details of the six specifications (1–6).

H Additional Counterfactual Results and Robustness

Table A11: Payment and Pricing Distortions Table, 2012

	Observed	Ban, fix p	Ban	Ban, $p = mc$	No Ban, $p = mc$
$Q_{statins}$	0.189	0.178	0.178	0.188	0.201
	(0.000)	(0.002)	(0.002)	(0.002)	(0.000)
$Q_{Lipitor}$	0.060	0.050	0.050	0.057	0.068
	(0.000)	(0.002)	(0.002)	(0.002)	(0.000)
$Q_{Crestor}$	0.026	0.018	0.018	0.025	0.035
	(0.000)	(0.002)	(0.002)	(0.002)	(0.000)
$OOP_{statins}$	16.11	16.11	15.80	2.66	2.66
	(0.0)	(0.0)	(0.5)	(0.0)	(0.0)
$OOP_{Lipitor}$	24.16	24.16	23.81	2.66	2.66
	(0.1)	(0.1)	(0.6)	(0.0)	(0.0)
$OOP_{Crestor}$	31.94	31.94	30.99	2.66	2.66
	(0.0)	(0.0)	(0.16)	(0.0)	(0.0)
$POS_{statins}$	57.60	57.60	56.33	188.17	214.57
	(0.1)	(0.1)	(0.21)	(6.25)	(5.86)
$POS_{Lipitor}$	67.39	67.39	66.35	332.67	375.46
	(0.1)	(0.1)	(0.16)	(12.18)	(12.25)
$POS_{Crestor}$	160.80	160.80	156.2	363.53	423.55
	(0.1)	(0.1)	(0.80)	(10.55)	(8.53)

Notes: Authors' calculations of observed and counterfactual equilibrium statin quantities (share of cardiovascular utilization) and prices (point-of-sale and out-of-pocket price per 30-day supply), per supply and demand model described in text. 2012 only. All simulations assume $\alpha^{de} = 0$, implying that prices and quantities are invariant to ε^{de}. "Ban, fix p" eliminates meals, holding POS and OOP prices fixed. "Ban" eliminates meals and allows both prices and quantities to adjust. "Ban, $p = mc$" eliminates meals and sets p^{oop} equal to marginal costs, then allows POS prices and quantities to adjust. Finally, "No Ban, $p = mc$" simply sets p^{oop} equal to marginal costs, then allows POS prices and quantities to adjust. Estimates based on 116,559 doctor-drug-brand-year observations with standard errors clustered at the doctor d level ($N_d = 13,793$) via delete-120 jackknife bootstrap.

Table A12: Welfare and Counterfactual Estimates – Supplement to Figure 6

	ε^{de}	-350	-100	-50	0
Total Surplus	Observed	90.21	44.75	35.66	26.56
		(0.63)	(0.64)	(0.64)	(0.64)
	Ban, 2011	-2.75	-0.46	-0.02	0.41
		(0.39)	(0.04)	(0.05)	(0.12)
	Obs, 2012	3.08	1.19	0.82	0.44
		(0.10)	(0.10)	(0.10)	(0.10)
	Ban, 2012	-0.35	0.56	0.77	0.93
		(0.58)	(0.09)	(0.06)	(0.07)
Consumer Surplus (retail)	Observed	88.33	42.87	33.77	24.68
		(0.63)	(0.64)	(0.64)	(0.64)
	Ban, 2011	-2.57	-0.29	0.13	0.57
		(0.31)	(0.05)	(0.12)	(0.22)
	Obs, 2012	3.45	1.56	1.18	0.80
		(0.10)	(0.10)	(0.10)	(0.10)
	Ban, 2012	0.20	1.13	1.30	1.44
		(0.50)	(0.07)	(0.07)	(0.15)
Consumer Surplus (-POS transfers)	Observed	80.28	34.82	25.72	16.63
		(0.63)	(0.64)	(0.64)	(0.64)
	Ban, 2011	-0.78	1.40	1.86	2.31
		(0.07)	(0.36)	(0.45)	(0.52)
	Obs, 2012	5.51	3.62	3.24	2.86
		(0.10)	(0.10)	(0.10)	(0.10)
	Ban, 2012	3.72	4.57	4.75	4.93
		(0.26)	(0.25)	(0.33)	(0.42)
Producer Surplus	Observed	9.93	9.93	9.93	9.93
		(0.00)	(0.00)	(0.00)	(0.00)
	Ban, 2011	-1.92	-1.92	-1.92	-1.92
		(0.38)	(0.38)	(0.38)	(0.38)
	Obs, 2012	-2.42	-2.42	-2.42	-2.42
		(0.00)	(0.00)	(0.00)	(0.00)
	Ban, 2012	-4.04	-4.04	-4.04	-4.04
		(0.35)	(0.35)	(0.35)	(0.35)

Notes: Authors' calculations of equilibrium surplus measures, in dollars per cardiovascular patient. "Meal Ban" counterfactuals allow both prices and quantities to adjust, per supply and demand model described in text. All simulations assume $\alpha^{de} = 0$; results shown for select values of ε^{de}. Estimates based on 116,559 doctor-drug-brand-year observations with standard errors clustered at the doctor d level ($N_d = 13,793$) via delete-120 jackknife bootstrap.

Table A13: Naive $\alpha^{de} = 0$ vs. Sophisticated $\alpha^{de} = 1$ Insurer Pricing (2011)

		"Naive" $\alpha^{de} = 0$ (main text)				"Sophisticated" $\alpha^{de} = 1$ (alternative)			
	ϵ_{BH}	-350	-100	-50	0	-350	-100	-50	0
Q Statins	Observed	0.18	0.18	0.18	0.18	0.18	0.18	0.18	0.18
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
	Ban	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Q Lipitor	Observed	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
	Ban	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Q Crestor	Observed	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
	Ban	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
OOP Statins	Observed	16.93	16.93	16.93	16.93	16.97	16.94	16.93	16.91
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)
	Ban	-0.48	-0.48	-0.48	-0.48	-0.05	0.13	0.22	0.38
		(0.06)	(0.06)	(0.06)	(0.06)	(0.03)	(0.07)	(0.10)	(0.14)
OOP Lipitor	Observed	32.35	32.35	32.35	32.35	32.42	32.36	32.33	32.30
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
	Ban	-0.84	-0.84	-0.84	-0.84	-0.12	0.22	0.39	0.71
		(0.08)	(0.08)	(0.08)	(0.08)	(0.07)	(0.15)	(0.21)	(0.30)
OOP Crestor	Observed	31.10	31.10	31.10	31.10	31.12	31.08	31.07	31.06
		(0.00)	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)
	Ban	-1.18	-1.18	-1.18	-1.18	-0.48	-0.07	0.12	0.45
		(0.18)	(0.18)	(0.18)	(0.18)	(0.05)	(0.07)	(0.09)	(0.18)
POS Statins	Observed	69.21	69.21	69.21	69.21	69.36	69.24	69.19	69.11
		(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)
	Ban	-2.03	-2.03	-2.03	-2.03	-0.35	0.41	0.79	1.46
		(0.25)	(0.25)	(0.25)	(0.25)	(0.08)	(0.27)	(0.38)	(0.57)
POS Lipitor	Observed	140.90	140.90	140.90	140.90	141.22	140.95	140.84	140.69
		(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.04)	(0.04)	(0.06)
	Ban	-3.66	-3.66	-3.66	-3.66	-0.53	0.94	1.72	3.08
		(0.34)	(0.34)	(0.34)	(0.34)	(0.31)	(0.66)	(0.89)	(1.29)
POS Crestor	Observed	133.75	133.75	133.75	133.75	133.84	133.69	133.64	133.56
		(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(0.03)	(0.04)
	Ban	-5.06	-5.06	-5.06	-5.06	-2.08	-0.29	0.52	1.92
		(0.77)	(0.77)	(0.77)	(0.77)	(0.21)	(0.30)	(0.40)	(0.77)

Notes: Authors' calculations of observed and counterfactual equilibrium statin quantities (share of cardiovascular utilization) and prices (point-of-sale and out-of-pocket price per 30-day supply), per supply and demand model described in text. 2011 only. Left panel: $\alpha^{de} = 0$, right panel: $\alpha^{de} = 1$. Results shown for select values of ε^{de}. Estimates based on 116,559 doctor-drug-brand-year observations with standard errors clustered at the doctor d level ($N_d = 13,793$) via delete-120 jackknife bootstrap.

Table A14: Robustness of Welfare Estimates to Modeling Assumptions

	ε^{de}	-350	-100	-50	0
Total Surplus, Ban 2011	Baseline	-2.75	-0.46	-0.02	0.41
		(0.39)	(0.04)	(0.05)	(0.12)
	Rebates	-2.66	-0.47	-0.04	0.37
		(0.38)	(0.04)	(0.05)	(0.11)
	MC	-2.76	-0.49	-0.04	0.39
		(0.39)	(0.04)	(0.05)	(0.12)
	Pricing	-2.90	-0.53	-0.06	0.39
		(0.42)	(0.04)	(0.05)	(0.12)
	ε^{de}	-7.59	-1.88	-0.74	0.41
		(3.41)	(0.90)	(0.42)	(0.12)
Consumer Surplus, Ban 2011 (Retail)	Baseline	-2.57	-0.29	0.13	0.57
		(0.31)	(0.05)	(0.12)	(0.22)
	Rebates	-2.43	-0.26	0.15	0.57
		(0.30)	(0.05)	(0.12)	(0.21)
	MC	-2.56	-0.29	0.14	0.57
		(0.31)	(0.05)	(0.12)	(0.22)
	Pricing	-2.78	-0.42	0.02	0.46
		(0.35)	(0.04)	(0.10)	(0.19)
	ε^{de}	-7.42	-1.71	-0.57	0.57
		(3.29)	(0.84)	(0.37)	(0.22)
Consumer Surplus, Ban 2011 (-POS Transfers)	Baseline	-0.78	1.40	1.86	2.31
		(0.07)	(0.36)	(0.45)	(0.52)
	Rebates	-0.51	1.58	2.01	2.44
		(0.06)	(0.39)	(0.46)	(0.54)
	MC	-0.77	1.41	1.87	2.32
		(0.07)	(0.36)	(0.45)	(0.53)
	Pricing	-1.14	1.09	1.53	1.95
		(0.10)	(0.30)	(0.37)	(0.45)
	ε^{de}	-5.72	-0.03	1.15	2.31
		(3.03)	(0.67)	(0.36)	(0.52)
Producer Surplus, Ban 2011	Baseline	-1.92	-1.92	-1.92	-1.92
		(0.38)	(0.38)	(0.38)	(0.38)
	Rebates	-2.10	-2.10	-2.10	-2.10
		(0.41)	(0.41)	(0.41)	(0.41)
	MC	-1.95	-1.95	-1.95	-1.95
		(0.38)	(0.38)	(0.38)	(0.38)
	Pricing	-1.73	-1.67	-1.63	-1.58
		(0.34)	(0.32)	(0.32)	(0.30)
	ε^{de}	-1.92	-1.92	-1.92	-1.92
		(0.38)	(0.38)	(0.38)	(0.38)

Notes: Authors' calculations of the effects of a meal ban on equilibrium surplus measures in 2011, in dollars per cardiovascular patient, for baseline specification (as in Figure 6 and Appendix Table A13), and alternative specifications: "Rebates" (alternative rebates as described in Appendix E); "Marginal Costs" (extreme alternative assumption that $mc = 0$); "Pricing" ($\alpha^{de} = 1$ rather than $\alpha^{de} = 0$); and an alternative specification with $\varepsilon_d^{de} = \gamma^{de} * \bar{\theta}_d^m$ rather than fixed ε^{de} across all physicians. Estimates based on 116,559 doctor-drug-brand-year observations with standard errors clustered at the doctor d level ($N_d = 13,793$) via delete-120 jackknife bootstrap.