

NBER WORKING PAPER SERIES

NO FREE LUNCH?
WELFARE ANALYSIS OF FIRMS SELLING
THROUGH EXPERT INTERMEDIARIES

Matthew Grennan
Kyle Myers
Ashley Swanson
Aaron Chatterji

Working Paper 24864
<http://www.nber.org/papers/w24864>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2018, Revised June 2022

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, Ben Handel, 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 constructed using data from The Nielsen Company, LLC and 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." The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Matthew Grennan, Kyle Myers, Ashley Swanson, and Aaron Chatterji. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

No Free Lunch? Welfare Analysis of Firms Selling Through Expert Intermediaries
Matthew Grennan, Kyle Myers, Ashley Swanson, and Aaron Chatterji
NBER Working Paper No. 24864
July 2018, Revised June 2022
JEL No. I1,L0

ABSTRACT

We study how firms target and influence expert intermediaries. In our empirical context, pharmaceutical manufacturers provide payments to physicians during promotional interactions. We develop an identification strategy based on plausibly exogenous variation in payments driven by differential exposure to spillovers from academic medical centers' conflict-of-interest policies. Using a detailed case study of an important class of cardiovascular drugs, we estimate heterogeneous effects of payments on prescribing, with firms targeting highly responsive physicians. Our model of supply and demand allows us to quantify how oligopoly prices reduce drug prescribing, and how payments move prescribing closer to the optimal level, but at great financial cost to patients and payers. In our estimated model, consumers are worse off with payments, unless there is substantial underprescribing due to behavioral or other frictions. In a final exercise, we calibrate such frictions using clinical data. We estimate that, in this case study, payments benefit consumers.

Matthew Grennan
University of California, Berkeley
Haas School of Business
Berkeley, CA 94720
and NBER
mgrennan@berkeley.edu

Ashley Swanson
Graduate School of Business
Columbia University
3022 Broadway, Uris Hall 621
New York, NY 10027
and NBER
ats2180@gsb.columbia.edu

Kyle Myers
442 Morgan Hall
Harvard Business School
Boston, MA 02163
kmyers@hbs.edu

Aaron Chatterji
The Fuqua School of Business
Duke University
100 Fuqua Drive, Box 90120
Durham, NC 27708
and NBER
ronnie@duke.edu

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 small but 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; Bhattacharya et al. 2020; David et al. 2010; Egan et al. 2019; Robles-Garcia 2020). In health care, physicians receive payments from pharmaceutical and medical device firms. The interactions that accompany those payments can 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 have long been a key component of drug promotions, and several studies have found a positive association between those expenditures and pharmaceutical prescribing.¹ The policy implications of such associations are difficult to interpret in light of several well-documented facts. First, physician treatment behavior varies widely (e.g., Cutler et al. 2019). Second, pharmaceutical firms spend large amounts of time and money targeting promotions to physicians; that is, payments are not allocated randomly (Fugh-Berman and Ahari 2007). Third, the equilibrium effects of payments depend on important features of the market, such as oligopolistic competition, insured demand, and other frictions—such as imperfect agency or behavioral biases—that might drive a wedge between the treatment a physician chooses and the treatment that maximizes patient welfare (Baicker et al. 2015; Besanko et al. 2020; Dickstein 2017; Inderst and Ottaviani 2012).

In this paper, we address these challenges using detailed data, a new instrumental variables (IV) strategy to estimate physicians' heterogeneous responses to payments from firms, and a flexible structural model of supply and demand that includes the role of payments as well as the possibility of choice frictions. We estimate wide variation in treatment effects across physicians, and we find that pharmaceutical firms target physicians with more positive expected responses to payments. We then use the estimated model, combined with a calibration exercise based on clinical trial results, to explore the equilibrium price, quantity,

¹See, e.g., Spurling et al. (2010) and Kremer et al. (2008) for reviews of early research on this topic. Other early research includes a marketing literature using data on "detailing" interactions for a subsample of physicians (Chintagunta and Manchanda 2004; Manchanda and Honka 2009; Narayanan and Manchanda 2009). More recent papers using data like those used here—which have payments for all physicians but do not enumerate other detailing interaction details—include Datta and Dave (2016); DeJong et al. (2016); Yeh et al. (2016).

and welfare impacts of a ban on payments.

This paper illustrates our approach using a detailed case study of the market for statins, one of only two large drug categories for which complete payment data was available in 2011-12. Statins are also interesting in and of themselves as one of the largest-selling drug categories in history. We discuss the tradeoffs involved in scaling to many drug categories in the Conclusion. We combine prescribing data from Medicare Part D with payment data at the physician-drug-year level and numerous other physician, hospital, and market characteristics related to payments and prescribing. The two branded statins in our sample, Pfizer's Lipitor and AstraZeneca's Crestor, were the first "strong statins" found to be differentiated from older generic statins in that they generated larger reductions in cholesterol. Lipitor and Crestor were heavily promoted drugs: over 75 percent of prescriptions in our data were written by a physician who received a meal from at least one firm (meals represent almost 98 percent of payment instances between statin firms and physicians). Lipitor and Crestor made up nearly 40 percent of statin prescribing in 2011 while maintaining prices around seven times those of generic alternatives. Thus, statins also provide a representative example of a market with firm payments promoting expensive branded drugs in conjunction with market power and other potential frictions.

We construct a flexible structural model as a lens through which to view this setting and a tool to simulate welfare and counterfactual equilibria. Firms negotiate prices with insurers and allocate sales force time and payments as a function of drug, regional, and physician characteristics. Patients visit physicians and fill prescriptions as a function of the drug's benefit to the patient, the out of pocket price the patient must pay, the effect of any payment interaction on physician decisions, and a decision error that captures the potential for physician recommendations to deviate from what is optimal for the patient.

Estimating the demand portion of the model proceeds in two steps. In the first step, we leverage Lipitor's patent expiration at the end of 2011. The ensuing generic entry generated a large shock to choice sets and relative prices: a new product with Lipitor's exact same molecule became available at a much lower price, and many insurers removed branded Lipitor from their formularies. This allows us to construct differences-in-differences style estimators for the parameters that determine price sensitivity and substitution across statins. The elasticity estimates that emerge are sensible and consistent with other estimates of pharmaceutical demand. This step also estimates a set of fixed effects for each physician-molecule, which encompass how drug quality, firm payment efforts, and any decision errors combine to determine physician prescribing conditional on price.

In the second step, we regress the physician-molecule fixed effects on payments, specifically meals. To obtain the causal effect of payments on prescribing, we use instrumental

variables based on variation in physicians' exposure to AMC's conflict of interest policies, which restrict firms' ability to provide payments to affiliated physicians. We consider payments to be an observable proxy for a variety of related interactions. In our main specifications, we use an inclusive version of this proxy: an indicator for a physician ever receiving a meal from a manufacturer in our data. Motivated by geographic economies of scale in firms' marketing efforts, we document that the effects of such policies spill over to *other* physicians who are unaffiliated with the AMC but happen to practice nearby.² The pattern of these spillovers 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 (Angrist et al. 2010).

Our IV approach diverges from several recent papers on physician-industry interactions that approach the issue of physician selection using specifications with physician fixed effects (Agha and Zeltzer 2019; Carey et al. 2020; Shapiro 2018a). The fixed effect approach is valuable for estimating certain treatment effects, but has limitations for this paper's goal of 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 larger than the within-physician effect of an incremental payment. Second, if firms target physicians based on their heterogeneous expected responses to payments, then the effect of a policy change on any measure of interest will depend on which physicians are treated in the baseline and counterfactual scenarios (Heckman et al. 2006).

Given these concerns, we focus on the cross-sectional variation in which physicians receive payments, and we use quasi-exogenous variation driven by physician exposure to AMC policies to estimate the distribution of marginal treatment effects (MTEs; Heckman and Vytlacil 2007) of payments on prescribing. In order to gain the power needed to precisely estimate MTEs, and also to control for physician- and market-level prescribing differences, we allow for flexible functions of our instruments and a large number of potentially relevant control variables. This creates a dimensionality and sparsity problem, which we address by drawing on the recent literature at the intersection of machine learning and econometrics. We use Lasso regressions to select the most powerful set of predictors and sample splitting

²This "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 promoted drugs at the institutions that impose them.

to ensure that our estimates are robust to errors in the variable selection process (Belloni et al. 2017; Chernozhukov et al. 2018).

This analysis yields an important new result: there is dramatic variation in physicians' responsiveness to payments. Our estimates imply that a meal payment relationship increases promoted statin prescribing by 41 percent for the average physician, which is roughly equivalent to the impact of a \$35 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 more positive expected treatment effects, (ii) would otherwise prescribe below-average shares 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.

To understand how the above demand parameters interact with other market frictions and ultimately influence welfare, we next analyze how prices are determined in equilibrium. We combine our demand estimates with a model of price negotiation between upstream manufacturers/distributors and insurers, and with external data on marginal costs, to capture the forces driving the point-of-sale prices that insurers pay for pharmaceuticals and, in turn, the out-of-pocket prices paid by patients. 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.

The final element needed to connect our model to welfare is a "decision error" parameter that captures the various reasons why physician decisions could be suboptimal for patients: variation in physician information and skill (e.g., Currie and MacLeod 2020), imperfect agency not driven by payments (Jacobson et al. 2006), or various behavioral biases (Baicker et al. 2015) are examples from prior research that may play a role here. In our model, payments could reinforce or counteract such frictions.³ 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 implied decision error in our estimated model.

Our counterfactual simulations yield several additional insights. First, the equilibrium effect of meals is to increase statin use by around five percent, and to increase use of the focal branded statins by 29 percent, on average. These equilibrium calculations are smaller

³In 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).

than the average treatment effects of meals on prescribing because they account for which physicians receive meals, and for the effects of business stealing in the case of physicians receiving meals from both firms. Second, high branded statin prices lead to prescribing below the efficient level in a world without meals. Our estimated model suggests that payments increase prescribing to near the efficient level (according to revealed preference demand estimates), though at high cost to consumers and payers.

Considering a range of potential decision errors, 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, decision errors bias revealed willingness to pay for statins downward by a substantial amount (if there is enough under-prescribing), then consumer surplus increases in the presence of meals. Whether the behavioral or other frictions underlying decision errors are severe enough to justify the allowance of meals is an empirical question and likely varies widely across contexts. For the case study of statins, we shed light on this question by calibrating a decision error value that fits the difference between our estimates of the average revealed preference willingness-to-pay and conservative estimates of the dollar value of life-year gains due to statins from the clinical literature. Our calibrated decision error value implies substantial under-prescribing of statins, well into the region where payments increase consumer welfare. Under this calibration, the implied 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 an important case study, this paper contributes a useful new instrumental variables strategy and a framework for estimating heterogeneous treatment effects of 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; Bhattacharya et al. 2020; Egan et al. 2019; Levitt and Syverson 2008; Schneider 2012) and in particular to the literature on drivers of physician treatment recommendations (Clemens and Gottlieb 2014; Dickstein 2017; Gruber and Owings 1996; Ho and Pakes 2014; Iizuka 2012). Our supply and demand model contributes to a nuanced literature on the nature of payments from firms to physicians in a range of important drug classes (Agha and Zeltzer 2019; Carey et al. 2020). Our focus on heterogeneity in treatment effects and targeted promotion adds new elements to a growing literature on the equilibrium effects of expert inducements in imperfectly competitive markets (e.g., Egan et al. 2020; Robles-Garcia 2020), and of direct-to-consumer drug advertising (see, e.g., Shapiro (2018b) and Sinkinson and Starc (2019)). 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. We build on prior work that has allowed advertising to be informative or persuasive, especially Dubois et al. (2018)'s study of junk food advertising. More broadly, our approach adds to a literature that has so far required unique data on which decision-makers are less subject to 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 presents our model of payments, pricing, and the demand for statins. Section 4 steps through our empirical models, identification and estimation approaches, and results. Section 5 presents the results of counterfactual simulations of a ban on physician-firm payments, including calibrating decision errors using clinical trial data. 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 describes institutional details of pharmaceutical markets in the US, and in particular the Medicare statin market in 2011-2012. It also describes our sources, sample restrictions, and summary statistics for data on (1) drug prices and quantities, (2) payments from firms to physicians, and (3) other regional variation in the data, with a particular focus on our research design that makes use of the Lipitor patent expiration as well as regional spillovers from AMC conflict-of-interest policies.

2.1 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) Statins are one of the few drug categories for which we observe payments from all branded manufacturers. 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) Though the data only specify the firm (not drug) associated with each payment, statins accounted for more than 80 percent of cardiologist-prescribed revenue for both Pfizer and AstraZeneca, making it likely that they are an important subject of any firm interactions with cardiologists.⁴ Also, although cardiologists accounted for only 10 percent of Part D statin claims, specialist prescriptions are often the first prescription written for a patient, which is then sustained by primary care physicians (PCPs) (Fugh-Berman and Ahari 2007). This gives specialists an outsized impact on total prescribing, and also suggests that much of the prescribing we document will be new prescriptions, where an active choice of drug is made. (4) Finally, Lipitor's patent expiration generated a large shock to statin prices and formularies, helping to identify other features of demand curves separately from payment effects.

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 branded and generic drugs. This sensitivity may be muted by various frictions, including enrollees' limited understanding of coverage and physicians' imperfect agency (Abaluck et al. 2018; Goldman et al. 2007; Chandra et al. 2010). Part D plan issuers' strategies and profits are 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 Prescribing: Prices and Quantities

We obtain data on physician specialties, affiliations, and demographics from the 2013 CMS Physician Compare database, which contains all physicians treating Medicare patients (CMS

⁴In contrast, in interactions with primary care providers (PCPs) in 2011-2, Pfizer might have promoted Celebrex, Enbrel, Lipitor, Lyrica, Norvasc, Plevnar, Pristiq, Viagra, and Zyvox, and AstraZeneca might have promoted Synagis, Toprol, Seroquel, Atacand, Nexium, Prilosec, and Symbicort. Consistent with this, after Lipitor went off patent, payments from Pfizer to cardiologists sharply declined, but no such trend break was observed for primary care physicians (PCPs). Results for all of our main analyses run on PCPs are available by request.

2013). Each physician's practice location is matched to one of 3,436 local Hospital Service Area (HSA) markets for hospital care and one of 306 Hospital Referral Region (HRR) markets for major tertiary care, according to the Dartmouth Atlas (CECS 2012).

Prescribing data are from publicly-available CMS Part D claims files for 2011 and 2012 (CMS 2012). These claims data describe total prescriptions (in 30-day supplies) and spending for each prescriber-drug-year. The data include prescribing 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, a hierarchy of drug categories that reflect similarities in drug mechanism and disease intended to treat. We focus on statin (ATC code = "10AAC") prescribing and use cardiovascular (ATC code = "C") prescribing as a proxy for the total number of patients seen by the cardiologist who might potentially need a statin in a given year.

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 10th 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. We restrict the sample to 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 average claim quantities and drug prices for our sample (we turn to variation across physicians below). On average, a physician in our sample writes about 3,800 Medicare prescriptions in the cardiovascular class per year, and roughly 700 of these are for statins. The effect of entry by generic atorvastatin in December 2011 is dramatic—in its first full year of availability, this new alternative accounted for more than 25 percent of cardiologists' statin claims.

The remaining columns of Table 1 summarize prices. In 2011, Lipitor and Crestor out-

Table 1: Prescribing Summary Statistics

	Prescription		Out-of-Pocket		Point-of-Sale	
	Count, mean		Price (\$), mean		Price (\$), mean	
	2011	2012	2011	2012	2011	2012
All Cardiovascular	3,602	4,156				
All Statins	626	733				
Crestor	90	102	31.86	31.85	137.09	160.33
Lipitor	138	37	32.04	62.62	139.48	163.92
Atorvastatin		189		9.67		32.45
Other Generics (3)	398	404	4.55	3.84	13.31	10.30

Notes: Based on a total of 89,754 cardiologist-drug-year observations, reporting statistics for those with non-zero use of any statin within a given year. Prescriptions (30-day equivalent) and prices derived from the Medicare Part D public use files. Out-of-pocket prices are plan enrollment-weighted averages of Part D enrollee cost-sharing per 30-day supply. Point-of-sale prices are plan enrollment-weighted averages of the total retail prices paid per 30-day supply when prescriptions are filled. One month is the modal supply per claim. See Appendices A.2 and B for details on variable and sample construction.

of-pocket (OOP) prices—the prices paid by the enrollee when filling a prescription—were about seven times those of generics. The full point-of-sale (POS) prices paid by insurers plus enrollees were three to four times OOP prices, and were similarly an order of magnitude higher for branded statins than 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 these features in our pricing model in Section 3.3 and explore the robustness of our assumptions 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 of other generics due to initially limited generic competition. Other generic drugs' prices also decreased slightly between 2011 and 2012. 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.⁵

⁵Branded 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

2.3 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 encourage use 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, interest had been growing for some time. By 2010, several states had begun to institute their own payment limitations and/or public reporting rules; a number of high-profile lawsuits required payment disclosure as a remedy (Guo et al. 2020); calls from politicians and patient advocacy groups were gaining momentum (Grassley 2009); and a number of firms, including Pfizer and AstraZeneca, began to publicly release comprehensive data on payments to physicians (Ornstein and Grochowski Jones 2015).⁶ These documents are the basis of our payments data, which were generously shared by Kyruus, Inc.⁷

Table 2 summarizes our data on payments from firms to physicians. As shown in Panel A, meals account for 98 percent of the payments we observe in our data.⁸ Panel B shows how the distribution of meal payments very closely maps the distribution of overall payments. The only exception is at the very top of the distribution, where a few physicians receive very large payments due to consulting, speaking, and travel fees or research grants. While this is an interesting group, we focus our analysis on meals since they are clearly the dominant form of payment in this setting.

Sixty-seven percent of physicians, representing 77 percent of cardiovascular prescriptions in our sample, received a meal from at least one of the branded statin manufacturers. Meal-based relationships are highly persistent over time: for the firm-years in our estimation be used by Medicare Part D enrollees, so we omit them from our analysis. In our supply side estimation in Section 4.3, we allow for higher rebates by Pfizer in 2012 and test robustness of our results to this choice.

⁶The 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, but was discussed for years prior to its implementation.

⁷The 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.

⁸Among physicians who received meals in our sample, 17 percent also had non-meal payments. Among physicians who did not receive meals, very few (0.16 percent) received non-meal payments.

Table 2: Payments Summary Statistics

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

	Any Kind		Meal		Travel/ Speak/Consult		Research	
	Claim		Claim		Claim		Claim	
	Raw	wgt.	Raw	wgt.	Raw	wgt.	Raw	wgt.
Crestor	0.615	0.770	0.607	0.761	0.015	0.027	0.001	0.001
Lipitor	0.338	0.443	0.317	0.417	0.014	0.027	0.001	0.001
Either	0.685	0.782	0.670	0.766	0.027	0.042	0.002	0.002

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

		mean	p10	p50	p90	<i>p99</i>
		Crestor	Any	432.7	15.0	58.5
	Meal	81.3	15.0	54.0	159.3	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	548.2	15.0	74.0	243.5	13,350.5
	Meal	97.9	15.0	65.0	203.5	596.5

Notes: Statistics calculated on 25,318 cardiologist-drug observations. Claim-weighted means use 2011 claims for weighting. In Panel (a), the "Either" category reports whether the cardiologist received payments from either firm. In Panel (b), which reports the distribution of total payments per cardiologist-drug-year (excluding zeroes), the "Either (+)" category reports the sum of payments across both firms.

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: the 90th percentile of the distribution of meal dollar values across (nonzero) observations at the physician-firm-year level was less than \$134. While these dollar values are small, they represent only a fraction of the total cost of the overall relationship (see Liu et al. 2020 and further discussion below), and research has shown that 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. This proxy for the physician-firm relationship is very inclusive, in that it is unlikely that any cardiologist in our sample has a significant relationship (detailing or otherwise) with one of the paying firms without ever receiving a meal. During these meals (and other interactions for which meals proxy), sales representatives target prescribers with drug information 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. 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, if we instead use an indicator for receipt of *any* type of payment (e.g., meals, consulting, speaking, travel, or research), and if we allow for spillovers across physicians in the same hospital or practice.

2.4 Regional Variation and Conflict-of-Interest Policies

2.4.1 Prescribing and Payments across the U.S.

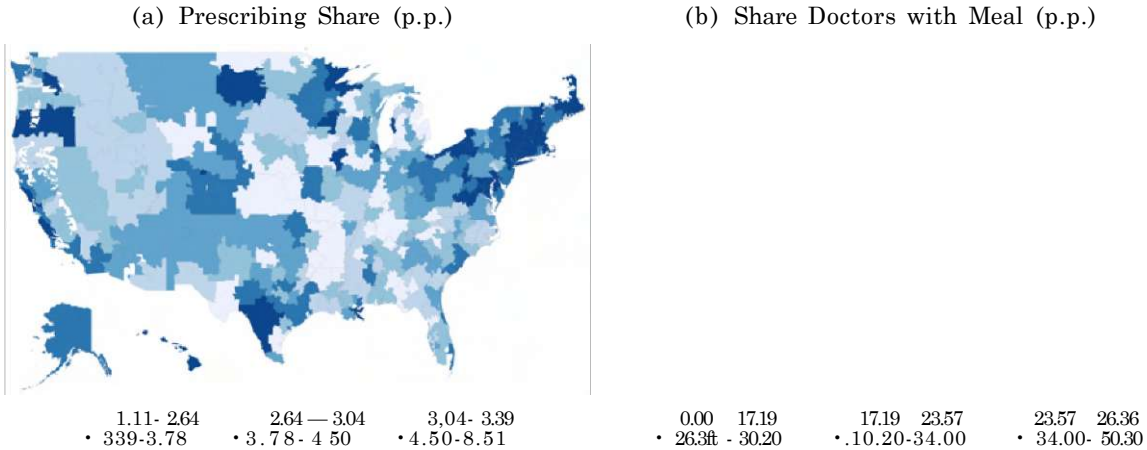
Figure 1 documents the geographic variation in utilization and meal payments across the U.S. Aggregating to the HRR level, Panel (a) plots the utilization of strong statins, and Panel (b) plots the share of cardiologists that receive meals from each branded drug manufacturer. Both show significant heterogeneity. For prescribing, cardiologists in the 90th percentile HRR were more than twice as likely than cardiologists in the 10th percentile HRR to prescribe a branded strong statin. Likewise, in the 10th percentile HRR, about 14 percent of cardiologists received a meal, while in the 90th percentile HRR about 36 percent of cardiologists received a meal (averaged across the two branded statin firms).

There is large geographic variation in both prescribing and payments, but 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 prescriptions written for Lipitor and Crestor, split by whether the physician received a payment from the focal firm. The two distributions are nearly identical in the raw data. If anything, there is slightly more prescribing of the focal drug among physicians who do not receive payments from its manufacturer.

2.4.2 Conflict-of-Interest Policies

To identify cardiologists who receive meals for plausibly exogenous reasons, we exploit the fact that, during the period we study, Academic Medical Centers across the U.S. had a wide range of policies 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 directly subject to them, but also 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

Figure 1: Regional Variation in Prescribing and Meal Payments, 2011



1.11- 2.64 2.64 — 3.04 3.04- 3.39 0.00 17.19 17.19 23.57 23.57 26.36
 • 3.39-3.78 • 3.78- 4.50 • 4.50-8.51 • 26.36 - 30.20 • 30.20-34.00 • 34.00- 50.30

(c) Prescribing Share (p.p.), by Meal Status

	mean	s.d.	p10	p50	p90
Share Prescribing Focal Drug (Meal)	3.71	2.58	1.17	3.12	6.96
(No Meal)	3.99	2.72	1.29	3.32	7.55

Notes: (a) the 2011 HRR-level averages of cardiologist-level cardiovascular shares for Crestor and Lipitor, averaging over both drugs. (b) the HRR-level share of cardiologists receiving meals from AstraZeneca or Pfizer, averaging over both firms. (c) the 2011 distribution of cardiologist-level cardiovascular shares for Crestor and Lipitor, averaging over both drugs, split by whether the same firm that produces the drug gave the cardiologist a meal. All numbers are in percentage points. Based on 25,318 doctor-drug-level observations from 2011.

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 characteristics of a geographic market for a given therapeutic area (Campbell 2008). 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). 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.⁹ Costs include the labor costs of additional sales representatives, the opportunity costs of diverting sales effort from

⁹For example, Alpert et al. (2019) document that Purdue Pharma avoided marketing OxyContin in states with strict prescription drug monitoring programs.

other physicians, and any direct costs of the interaction (e.g., meal expenditure). They also implicitly include factors that limit or prohibit access for sales representatives. For example, ZS Associates *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).

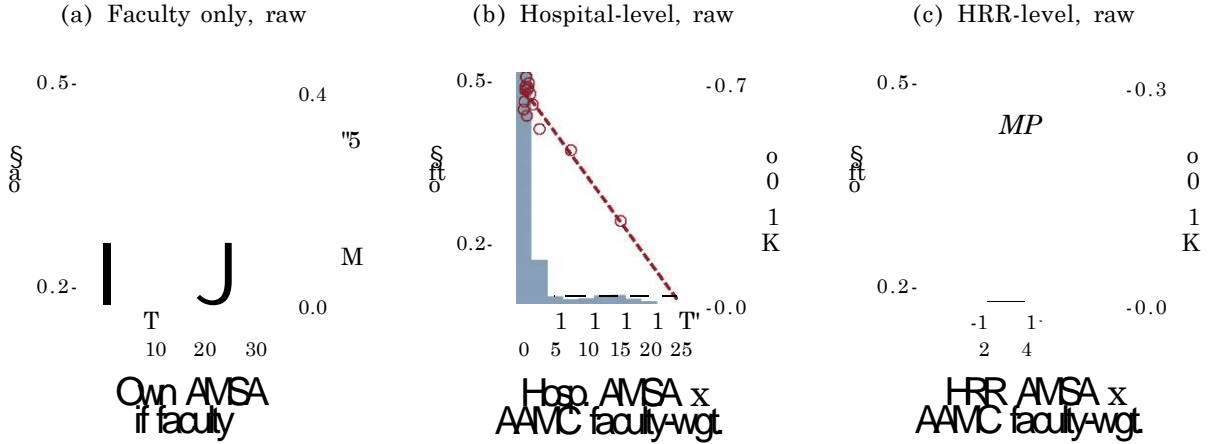
To operationalize these ideas, we link physicians to AMCs using the Association of American Medical Colleges (AAMC) faculty roster, and we obtain data on AMCs' conflict of interest policies from the American Medical Student Association's (AMSA) conflict of interest scorecard. The AMSA scores, ranging from 11 to 32 in 2011-12, 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.¹⁰ 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, and even to 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.

The binned scatterplots in Figure 2 illustrate the relationships between meal receipt and different measures of AMSA CoI scores. Faculty at AMCs with more stringent policies are less likely to receive meal payments (Panel (a)). This phenomenon is also observed for non-faculty physicians working at the same hospitals as faculty (Panel (b)). It also spills over at the regional level (at the HRR level shown here, and at the HSA level as shown in Appendix G.2)—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 pharmaceutical sales executives and pharmaceutical marketing consultants regarding economies of scale in sales force allocation.¹¹

¹⁰In 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 by summing across all AMSA components. See Larkin et al. (2017) for more details on the scorecard.

¹¹Personal communications: George Chressanthos Jan 5, 2018 and Pratap Khedkar Feb 15, 2018.

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



2.4.3 Physician-Level and Regional Characteristics

The primary concern with using these CoI 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). To help ensure that our identifying variation is driven by spillovers from CoI policies rather than these other factors (e.g., preferences, market structure, etc.), and to provide statistical power, we control for a rich set of observable physician and regional characteristics. Here, we outline how we use these controls in our research design below and provide an overview of the data.

From the CMS Physician Compare data, we observe each cardiologist's gender, year of medical school graduation, faculty status, the numbers of different organizations and practice locations listed as affiliations, and whether the physician is enrolled in the CMS's programs for electronic prescribing, electronic health records, or quality reporting.

We supplement this set of cardiologist-specific characteristics 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. Appendix G.1 reports the summary statistics for all controls and instruments, along with the results from univariate regressions of our utilization and meal payment variables on each control and instrument.

3 A Model of Payments, Pricing, and Demand for Statins

This Section presents a flexible structural model, motivated by the above institutional details and economic theory, that we use to estimate demand for statins (in particular the causal effect of payments on demand) and quantify welfare under the status quo as well as counterfactual scenarios where payments are banned. In our model, insurers negotiate point of sale prices with upstream suppliers, manufacturer sales representatives target meals to physicians, and physicians prescribe drugs. Because prices and payments depend upon expected demand, our discussion begins there.

3.1 Demand with Payments and Decision Errors

This Section develops an explicit model of how physicians and patients trade off the influences of meals and out-of-pocket prices and substitute across competing drugs, allowing for potential "decision errors" that drive a wedge between prescribing decisions and true patient utility. Let the indirect decision utility of drug $j \in J = \{0, 1, \dots, J\}$, for use case i (a doctor/patient/visit combination) in each market defined by doctor d in year t , be given by: $U_{idjt} = S_{dj} + \beta_{idjt}$.¹² The choice $j = 0$ represents the choice of treatment other than a statin, with mean utility normalized to $S_{d0t} = 0$. We measure the market size of potential statin patients for each physician-year as the number of all cardiovascular prescriptions, as a proxy for the number of patients who might potentially need a statin. The use-specific i.i.d. unobservable $\epsilon_{idjt} = \epsilon_{idt} + (1 - \alpha)\epsilon_{idjt}$ is the random coefficients representation of the nested logit model (Cardell 1997), where ϵ_{idt} is a random 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 drugs. As the nesting parameter $\alpha \in [0, 1]$ approaches

¹²The 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.

1, there is less substitution outside the nest of statins.¹³

We specify mean utility across use cases as:

$$u_{djt} = \mathbb{1}_{\{m_{dj} > 0\}} \cdot \theta_{pdjt} + x_{djt} \beta_j + \epsilon_{djt} . \quad (1)$$

Here, $\mathbb{1}_{\{m_{dj} > 0\}}$ is an indicator for whether cardiologist d received a meal from j 's manufacturer and its utility weight. Importantly, this utility weight may be specific to the drug-doctor pair, with arbitrary correlation patterns. It may even be negative and lead to decreased prescribing (e.g., due to new information received during the interaction accompanying the payment). While we are not able to micro-found the mechanisms underlying this heterogeneity, it likely captures the net effects of several sources of variation that have been discussed in prior research (e.g., Inderst and Ottaviani 2012), such as: 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, θ_{pdjt} is the average out-of-pocket price paid by patients and its weight. X_{djt} is a rich set of covariates that captures perceived quality variation across drugs, as well as regional and cardiologist variation in prescribing patterns over time (we discuss this in detail when we turn to estimation of the model in Section 4.1). Finally, ϵ_{djt} is a cardiologist-drug-year-level unobservable term, which we allow to have two components:

$$\epsilon_{djt} = idjt + e_{jt} . \quad (2)$$

ϵ_{djt} is a typical demand unobservable that impacts both choices and true realized utility. e_{jt} is a "decision error" in the spirit of Baicker et al. (2015) that affects consumer decisions but does not affect consumer surplus directly.

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). In Section 5, 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. At this stage, we leave the decision error specification fully flexible in terms of the mean decision error, heterogeneity in errors

¹³In Appendix G.4, we show the results of alternative specifications without a statin nest, and with a two-level nesting structure with a statin nest and another nest just for "strong statins."

across physicians and drugs, and the correlation with meal payment effects.

Given a set of drugs J_{dt} available to a cardiologist and flow of choice opportunities Q_{dt} , we assume the cardiologist/patient chooses the drug that maximizes decision utility, so that expected quantities demanded are given by:

$$Q_{djt} = Q_{dt} \Pr\{u_{iajt} > U_{ikdt}, \forall k \in J_{dt}\} = Q_{dt} \frac{p_j^i \exp(\beta \Delta u_{iajt})}{\sum_{k \in J_{dt}} p_k^i \exp(\beta \Delta u_{ikdt})} \exp(\beta \Delta u_{iajt})$$

Given this model, we represent expected consumer surplus as:

$$CS_{J} = Q_{dt} \ln \left(\sum_{j \in J_{dt}} p_j^i \exp(\beta \Delta u_{iajt}) \right) - \sum_{j \in J_{dt}} p_j^i \exp(\beta \Delta u_{iajt}) \Delta u_{iajt}$$

CS implied by decision utility adjustment for "decision errors" and meals

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- ($e_j < 0$) or over-prescribing ($e_j > 0$), as well as the countervailing (or reinforcing) effect of meals.¹⁴¹⁵ See Appendix C.2 for further discussion.

3.2 Targeting Meal Payments to Physicians

The parameter β in the demand model 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. Meal decisions are likely based on data we have available as researchers, plus other factors that are unobservable to us. We capture this by specifying a selection equation that is a semi-parametric representation of a model of strategic meal

¹⁴A related (and not mutually exclusive) interpretation would be that physicians maximize a sum of physician (chooser) and patient (consumer) utility, with β governing the difference between the physician's maximum and the patient's.

¹⁵In their study of banning advertising for junk food, Dubois et al. (2018) also allow decision utility to diverge from welfare relevant utility, considering cases where advertising affects decisions but not utility or enters utility directly. In our model, payment effects never enter welfare directly per se, but they can be arbitrarily correlated with welfare improvements, depending on the correlation between payments and any underlying decision errors.

allocation. This first stage selection equation takes the form of a linear probability model:

$$1[m_{dj} > 0] = Z_{dj}Y + X_{dj}Y + \hat{\alpha}_{dj} . \quad (5)$$

Appendix C.3.1 shows the tight relationship between Equation (5) and a structural version of this model for a particular cost function with increasing returns to scale.

3.3 Pricing Pharmaceuticals

The details of pharmaceutical supply chains are notoriously complicated. We seek to abstract from less relevant (for our purposes) details while capturing enough of the key economics of pharmaceutical pricing to generate credible estimates of the direction and magnitude of equilibrium price changes under 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).¹⁶

Let the supplier's profit be: $\pi_{djt} = p_{jt} - mc_{jt}$, where p_{jt} is the manufacturer rebate and mc_{jt} captures the cost of manufacturing and distributing the marginal unit of drug j . p_{jt} is the point-of-sale price insurers pay for the drug, which we model as constant across cardiologists within region r . We link the negotiated point-of-sale price and out-of-pocket price paid by enrollees via $p_{jt} = cs_{djt}P^oS$, where cs_{djt} is a cost-sharing parameter that varies across markets and years, depending on drug mix and insurer mix (discussed in detail in Appendix A.2). This reflects the practice whereby 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 physician's 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.¹⁷

We assume that prices of substitute drugs in the market are determined in a simultaneous Nash Equilibrium of Nash Bargaining between suppliers and buyers (Crawford and Yurukoglu 2012; Collard-Wexler et al. 2017). In the model, each price maximizes the Nash

¹⁶As discussed by Starc and Swanson (2020), both pharmacies and pharmaceutical manufacturers have market power, but relative market power of different suppliers varies by drug. These details are captured in a reduced form sense by the bargaining and cost-sharing parameters in our model below, which will be held fixed in our counterfactual analyses. This approach implicitly 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.

¹⁷Standalone Part D plans enter and negotiate prices in one of 34 Part D pricing regions, which 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.

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:

$$P_j K^i \cdot j = m^c j + j^t \left(1 + \frac{\xi}{\text{der}} \frac{dq_{j,t}}{dp_j} \right) \frac{E_{\text{der}}}{Y^{\wedge} \text{der}} \frac{CS_{dt}(J_{dt}) - CS_{dt}(J_{dt} \setminus j)}{Y_{\text{der}} \cdot q_{j,t}} + j^t (1 - j - m^c j)$$

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 drug j : $CS_{dt}(J_{dt}) - CS_{dt}(J_{dt} \setminus j)$. 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} . The CS function represents surplus from the insurer's perspective and thus differs slightly from CS as defined in Equation (4). We follow recent papers on insurer-hospital bargaining (Gowrisankaran et al. 2015; Ho and Lee 2017) by using a parameter $\alpha^{cs} \in [0,1]$ to capture the relative weight insurers place on enrollee surplus and plan costs:¹⁸

$$CS_{dt}(J_{dt}) := \alpha^{cs} Q_{dt} \ln \left(1 + \left(\frac{dq_{j,t}}{dp_j} \right) \frac{Y^{\wedge} \cdot q_{j,t}(p^{oop}) - j - p_j}{Y_{\text{der}} \cdot q_{j,t}} \right) + (1 - \alpha^{cs}) \frac{dq_{j,t}}{dp_j} \frac{Y^{\wedge} \cdot q_{j,t}(p^{oop}) - j - p_j}{Y_{\text{der}} \cdot q_{j,t}}$$

CS implied by decision utility insurer drug costs

We assume that insurers negotiate drug prices as a function of consumer surplus as implied by decision utility; intuitively, insurers suppose "doctors know best" when negotiating prices. Appendix C.5 relaxes this assumption and Appendix H shows that our qualitative findings are unchanged even in the opposite extreme where insurers perfectly adjust consumer surplus for decision errors and meals.

4 Demand and Supply Estimation and Results

In this Section, we show how meal payments can be fit into a standard potential outcomes framework, integrated with the structural demand system. The primary results are the

¹⁸In contrast to these papers, we model pricing of drugs within a single product class (statins), rather than pricing of a large bundle of products. Thus, α^{cs} in our setting may also capture how plan enrollment would respond to disagreement in this particular product class.

demand parameter estimates, with a particular focus on the heterogeneous treatment effects of meals on prescribing. The Section concludes with estimating the pricing model that is used for computing counterfactual equilibrium prices in the next section of the paper.

4.1 Demand Identification and Estimation

Our demand estimation approach proceeds in two broad steps. We outline the strategy here and describe each step in more detail in the remainder of the Section.¹⁹ In the first step, we estimate the price and nest parameters and a set of drug-doctor fixed effects, instrumenting to account for the endogeneity of prices and nesting patterns. In the second step, we set up a potential outcomes framework where the drug-doctor fixed effects are the outcome of interest and the key endogenous variable is the indicator for meal payments. Within this framework, we use our policy-spillover instruments to estimate the distribution of treatment effects across the sample of drug-doctor pairs.

We linearize the demand model, following Berry (1994). We set choice probabilities implied by the demand model in Equation (3) equal to observed market shares, and invert the system of equations to obtain mean utilities as a function of the market shares: $\delta_{dj_t} = \ln(s_{dj_t}/s_{dot}) - \ln(s_{dj}|_{gt})$. Combining this with Equation (1) yields the linear specification:

$$\ln(s_{dj_t}/s_{dot}) = \alpha \ln(s_{dj}|_{gt}) - \beta_j + \gamma_j + \delta_j + \epsilon_{dj_t} \quad (7)$$

where s_{dj_t} is j 's overall market share, s_{dot} is the market share of the outside good (non-statin treatments), and $s_{dj}|_{gt}$ is j 's market share within nest g , the set of statin treatments. During 2011-2012, non-statin treatments included lifestyle changes such as dietary modification and exercise, and several pharmaceuticals with less cholesterol-reducing efficacy than statins such as ezetimibe, bile acid resins, niacin, and fibrates (Harvard Men's Health Watch, 2014).

The following provides the empirical models and overview of the estimation routine, with the specifics detailed in Appendix D.

4.1.1 Estimating Price and Nest Parameters

In the first stage of estimation, we implement a differences-in-differences style estimator, leveraging the price and choice set variation resulting from the introduction of generic atorvastatin at the end of 2011 to identify the coefficients on price and within-nest share. We estimate:

$$\ln(s_{dj_t}/s_{dot}) = \alpha \ln(s_{jt}) - \beta_j + \gamma_j + \delta_j + \epsilon_{dj_t} + \eta_{Lip12} + \epsilon_{Cdj_t} \quad (8)$$

¹⁹Appendix D provides full step-by-step details on our algorithm.

where t^{aj} is a drug-doctor-specific fixed effect reflecting heterogeneity in doctors' mean decision utility over different treatments and $6t$ is a year fixed effect reflecting the (possibly evolving) average preferences over statins vs. the composite outside good of non-statin treatments over time. We further include 6_{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.

We account for the endogeneity of $\ln(s_{dj-|91})$ and p_j by constructing instrumental variables that 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).²⁰ 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, resulting in large variation in the relative prices and choice sets consumers faced. To utilize this variation, we create instruments for each plan-drug-year-region as the average out-of-pocket price for that insurer-drug-year across *other* regions. We then average across plans, weighting by enrollment, to create an instrument for physician d 's region. We also create an analogous instrument based on an average dummy for formulary inclusion. The instrument set is then: $Z^p = [p^{0j}P^{IV}, I_j^{formiv}]$. These are similar in spirit to the bargaining ability instruments in Grennan (2013, 2014) and Dickstein (2017). 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 $Z^j = [\ln(|Jdt|), |Jdt|, |Jdt|^2, \ln(|Jd?|), |Jd?|, |Jd?|^2]$ as instruments. This leverages the fact that more variety will mechanically affect within-group shares.

4.1.2 Estimating the Effects of Meals on Prescribing

The fixed effects $\hat{\alpha}_{dj}$ from the first step of our estimation capture all of the sources of persistent prescribing differences across doctors during our sample period. To estimate the extent to which these are influenced by meal payments from pharmaceutical firms, we project the drug-doctor fixed effects on our cross-sectional meal indicator and a rich set of controls for physician and market characteristics:

$$4 = 6mmi\{mdj>0\} + 6j + Xdj6^* + \text{\$}dj . \quad (9)$$

²⁰An additional challenge is that we observe average out-of-pocket prices at the drug-year-region level, implying that there is measurement error. Under the assumption that this is classical measurement error, our instruments for out-of-pocket price, which are primarily intended to address the endogeneity of price, will address this source of bias.

The idea of a secondary regression to uncover the determinants of fixed effects goes back at least to Mundlak (1978). To account for the fact that the fixed effects are measured with noise, we employ a version of the standard shrinkage approaches from the empirical Bayes literature.²¹ In our preferred specification, we construct $I_{\{mdj>0\}}$ 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 Intuitively, this approach aims to recover the steady-state effect of meal payments on prescribing. We estimate this equation only on observations for Lipitor/atorvastatin and Crestor, as generic firms do not provide meals to doctors.

The outcome equation (9) and selection equation (5) fit into the canonical potential outcomes framework. In the context of the model, the unobservable in the selection equation Δ_{dj} may be correlated with both Δ_{dj} and the heterogeneous component of θ_{jj} . In such a case, the standard 2SLS estimator will be a particular weighted average of the local average treatment effects on compliers, and this weighting may be of limited relevance for our policy simulations. We thus estimate the marginal treatment effects directly, using the `mtefe` package in Stata 16 (Andresen 2018). We can then estimate many treatment effects of interest as a function of these MTEs.

The cross-sectional nature of our identification strategy and the data-intensive nature of our semi-parametric MTE estimation make a rich set of controls 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 Section 2.4.3. Our instruments are built by starting with the aggregate AMSA scores of AMCs in the same region (HSA or HRR) as the focal cardiologist, excluding AMCs affiliated with the focal cardiologist, or the hospital where she works. We also interact each of these aggregates with variables that capture the intensity of spillovers at the individual cardiologist level: the percent of local cardiologists affiliated with the AMC and the drive time from the cardiologist's office to the AMC. Finally, we include linear, quadratic, and logarithm transformations of each instrument and control variable, and further interact each instrument with dummy variables for each manufacturer. This rich specification introduces issues with sparsity and collinearity that 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 regressions to select the controls and instruments which most strongly predict meals and

²¹ See Chandra et al. (2016) for a recent application in the health care 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 distribution.

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

prescribing.

4.1.3 Estimation Routine Overview and Inference

After obtaining the point estimates for the price and nest parameters using the full sample, the remainder of our estimation and inference routine is performed using 500 bootstrap iterations. Within each iteration, we first drop a random sample of vNd cardiologists and resample each remaining cardiologist's prescribing choices. We then estimate price and nest parameters for that sample; our reported standard errors for those parameters are the standard deviations of the 500 point estimates. At this point, we shrink the physician-molecule fixed effects toward the mean to account for potential measurement error. For the MTE estimation, we follow Chernozhukov et al. (2018) by splitting each bootstrap sample into two subsamples, separately estimating the Lasso and MTE models on opposite halves of the data, and taking the median of those two estimates. Our reported point estimates and standard errors are the median and median deviation of the resulting 500 estimates. Appendix D presents the estimation routine in full detail, and also presents results on the most frequently selected instruments and controls. For the supply estimation and counterfactuals, we again follow the median-based approach of Chernozhukov et al. (2018), since these estimates are based on results from the Lassos. The variation in drug prices is across states over time, so we jackknife a random set of seven (v^{50}) states in each of the 500 bootstrap iterations for the purposes of constructing the standard errors.

4.2 Demand Parameter Results

4.2.1 Price Coefficients and Substitution Patterns

Table 3 provides details on the estimates of the first step of the demand model, illustrating the importance of the rich fixed effects and instrumental variables in obtaining these results. (Appendix G.4 provides further details, including the robustness of these results to alternative specifications.) Focusing on our main model in the final column, the nesting parameter estimate of 0.42 is consistent with the 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 $n^p = fpp$ 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). 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'

Table 3: Demand Estimates Step 1—Price and Nest Coefficients and

	OLS, V_d	OLS, \hat{d}_j	IV, \hat{d}_j
Q^P	0.00106 (0.00002)	-0.00025 (0.00002)	-0.00753 (0.00016)
A	0.940 (0.0004)	0.966 (0.001)	0.423 (0.010)
mean(n^P)	0.262	-0.1098	-0.210
s.d.(n^P)	0.296	0.124	0.226
N	117,517	117,517	117,517
F-stat.			476.9
$R^2(\delta_{ajt} : \hat{d}_j - Q^P)$			0.809
mean($\hat{d}_j / 1 Q^P$) strong statins			-296.7
mean($\hat{d}_j / 1 Q^P$) other generics			-315.7
s.d.($V_d / Q^P $)			74.6

Notes: Reports parameter estimates from Eq. 8. Standard errors for the main parameters (Q^P and A), in parentheses, are based on the standard deviation of the point estimates from the 500 perturbed-bootstrap samples.

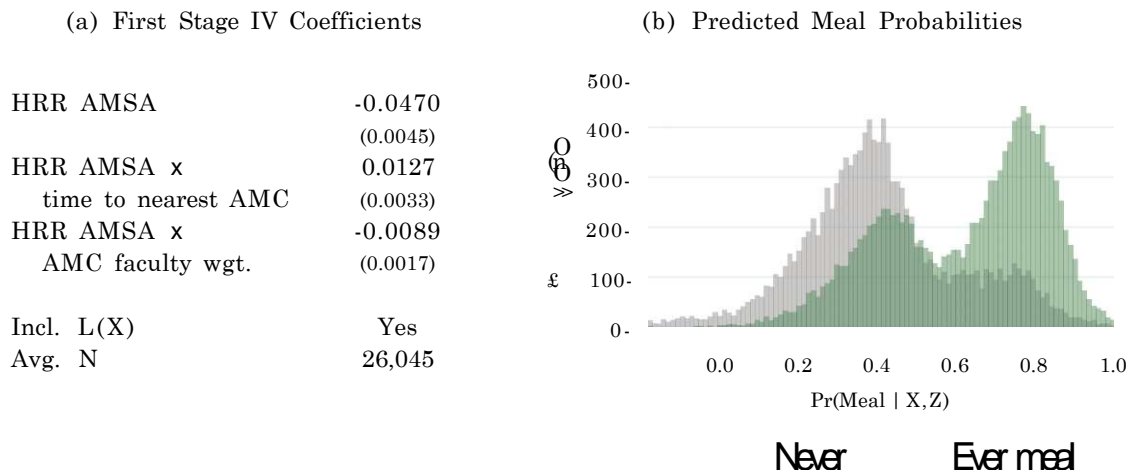
OOP). The overall physician-molecule preference variation itself is large, with one standard deviation of the distribution equivalent to an OOP price differential of about \$75.

4.2.2 Meal Payments First Stage

To explore the first stage effects of our AMSA instruments on meal payments, we first regress the meal payment indicator on the set of Lasso-selected controls and three intuitive transformations of the regional AMSA Col scores: (1) the average Col score for the AMCs in a region (excluding AMCs affiliated with a focal cardiologist's hospital), (2) the same score multiplied by the share of all doctors in the region that are faculty, and (3) the same score multiplied by driving time from a cardiologist's primary practice location to the nearest AMC. Figure 3 Panel (a) reports the results of this regression using the HRR-level IVs (see Appendix G.2 for similar results using the HSA-level IVs). We find that practicing in regions with AMCs with higher AMSA scores is negatively associated with receiving a meal; this effect is weaker if the focal physician's office is located farther from the AMC, and stronger if a larger percentage of doctors in the region are associated with the AMC. These patterns are consistent with our access costs and spillover hypotheses, and we find that they hold even when cardiologists who are faculty themselves are excluded from the sample (see Appendix G.6). Appendix D.3 provides more details on the most frequently selected controls and instruments from our full estimation routine. Across our bootstrap iterations, we estimate a median cluster-robust weak identification F-statistic of 136.7 (median deviation=40.4). However, in settings like ours, where we use Lasso to select from a high-dimensional set of covariates, the traditional statistical tests for IV relevance have been shown to perform

poorly. Thus, we also estimate versions of the sup-score weak identification tests proposed by Belloni et al. (2012) and implemented within the `ivlasso` Stata package (Ahrens et al. 2020): all bootstrap subsamples are able to reject the null hypothesis of weak instruments at a critical value of 0.05.

Figure 3: Explaining Variation in Meal Payments



Notes: (a) reports the results from an OLS regressions of a meal indicator on the vector of Lasso-selected controls, $L(X)$, and manually selected HRR-level instruments, which are standardized so that the coefficient indicates the percentage point change in predicted meal probability given a one s.d. change in the IV (b) displays the distribution of predicted meal probabilities, split by actual treatment status, from the full estimation routine where both controls and IVs are selected by Lasso.

Our identification strategy requires that Col scores be powerful predictors of meal payments, and that they only affect prescribing through the channel of meal payments. While we cannot test the exclusion restriction directly, we explore its validity by conducting a set of placebo tests in the spirit of Angrist et al. (2010). In our context, there are subsamples of physicians whose observable characteristics make them very unlikely to receive a meal payment, such that the causal relationship between the IVs and meals is substantially shut down. 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 significant first stage and reduced form relationships disappear, as we would expect if the exclusion restriction were satisfied.

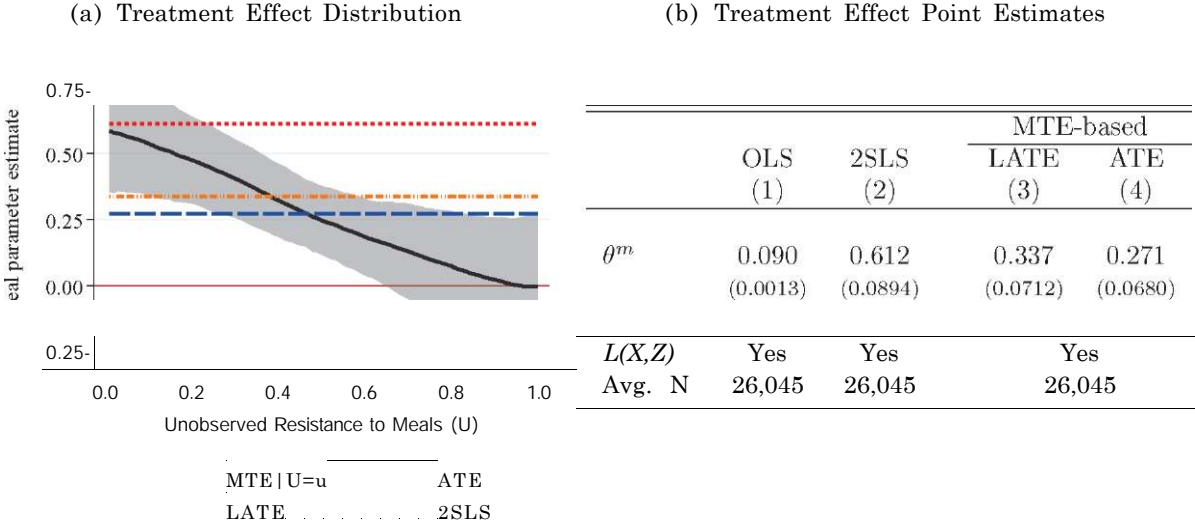
Figure 3, Panel (b) shows histograms of the first stage propensity score estimates (predicted meal probabilities) from the full estimation routine, for physician-drug observations with and without meal payments. The model produces large overlapping support for the two groups across the unit interval. In our application, we have found that the rich specifica-

tion 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 importance of a rich specification for the first stage.

4.2.3 Marginal Treatment Effects of Meals on Prescribing

Figure 4, Panel (a) plots our MTE estimates vs. the unobserved resistance to treatment.²³ The average treatment effect of 0.27 is roughly equivalent to the effect of a \$35 price decline, but we reject the hypothesis of a homogenous treatment effect. At the 10th percentile of unobserved resistance (i.e., physicians that firms are very likely to pay), the effect is 0.53 (equivalent to a \$71 price decline or a 1.0 standard deviation increase in the underlying physician-molecule preference heterogeneity), while at the 90th percentile of unobserved resistance (i.e., physicians that firms appear to avoid), the point estimate is 0.02 (a \$3 price decrease or 0.04 standard deviations of the underlying heterogeneity) and it is not statistically distinguishable from zero. Appendix G.6 shows that the level and slope of these MTE estimates are similar under alternative samples and modeling decisions.

Figure 4: Marginal Treatment Effects of Meal Payments



Notes: (a) plots the MTE curve ($E[\theta^m | U = u]$) with 95 percent C.I. in shaded grey, alongside other point estimates of θ^m . (b) displays the estimates from the OLS (Col. 1), 2SLS (Col. 2) and MTE (Cols. 3-4) specifications, all of which include the Lasso-selected instruments and controls, $L(X,Z)$.

The table in Panel (b) compares several estimates of θ^m : ordinary least squares (column

²³The 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.

(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 Heckman et al. (2006) suggest for cases with positive selection on heterogeneous treatment effects, the 2SLS estimator overestimates the ATE obtained from the MTEs—that is, 2SLS would overestimate the effect of banning payments by a factor of two. Our estimated average treatment effects are still 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,²⁴ whereas the coefficients in our nested logit demand model imply that a meal-based relationship increases promoted statin prescribing by about 40 percent for the average physician, but by 60 percent for the average physician actually targeted by firms and only 20 percent for physicians firms avoid. This could be due to differences between statins and other drug categories, or because the effect of the overall relationship may be much larger than the within-physician effect of an incremental meal.²⁵ 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.

An advantage of the MTE estimation approach is that the resulting estimates can be paired with the data (i.e., physician observables and realized treatment) to derive the expected response to treatment $E[\theta^j]$ for any observation in the data.²⁶ Figure 5, Panel (a) presents a histogram of expected treatment effects, normalized by the standard deviation of the physician-molecule preference variation, $E[\theta^j]/SD(\theta^m)$, separately for physicians with and without meal payments. Payments are clearly directed to physicians with more positive expected responses to treatment.

The median expected response of those receiving payments is roughly a 0.71 standard deviation change in the mean preference for prescribing the focal drug. By contrast, for those not receiving payments, the analogous median effect is roughly 0.24, and for roughly 80 percent 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 4 Panel (b), which implied a sizable

²⁴We credit Carey et al. (2020) for this calculation.

²⁵Chintagunta 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.


²⁶More formally, $E[\theta^j | X_{aj}, I\{m_{aj} > 0\}]$; see Appendix D.2, and Eq. 16 specifically, for more on how individual-level expected treatment effects are derived from the MTE model.

difference between the average treatment effect on the treated and the average treatment effect on the untreated.

Figure 5: Heterogeneity in Expected Treatment Effects Across Doctors, by Actual Treatment Status


(a) Prescribing Change with Payment (SD)

(b) Prescribing Without Payments (share)



	p10	p25	p50	p75	p90		p10	p25	p50	p75	p90
Never	-0.004	0.118	0.241	0.352	0.458	Never	0.013	0.021	0.033	0.052	0.076
	(0.176)	(0.174)	(0.175)	(0.174)	(0.171)		(.)	(.)	(.)	(.)	(.)
Ever	0.485	0.601	0.706	0.812	0.908	Ever	0.006	0.010	0.016	0.026	0.040
meal	(0.146)	(0.150)	(0.151)	(0.148)	(0.149)	meal	(0.001)	(0.001)	(0.002)	(0.003)	(0.005)

(c) Manufacturer Return to Payment (\$)



	p10	p25	p50	p75	p90
No meal	-1688.5	-1244.9	-441.5	1,085.2	3,447.5
	(620.9)	(679.9)	(1,044.6)	(1,813.5)	(2,851.1)
Meal	-731.7	11.3	1,623.7	4,551.6	8,739.9
	(187.4)	(304.8)	(550.6)	(962.1)	(1,519.5)

Notes: Plots the posterior (dj-) cardiologist-drug-specific estimates of: (a) expected meal response parameters $E[0d_{ij}]$, scaled by the standard distribution of the cardiologist-drug mean utility "fixed effects" r^{dj} ; (b) the distribution of expected cardiovascular shares, setting each meal indicator to zero $E[s^{TM=0}]$; and (c) the distribution of manufacturers' incremental profits due to meals $(p^{0t} (1 - T_{jt} - w_{jt}) - m^{crt}) \times E[q^{ij=1} - q^{ij=0}] - C_j$, given the price (p^{pos}) paid by insurers and consumers, net of marginal production and distribution costs mc , costs of payment relationships C_j , rebates T , and markups charged by supply intermediaries w (see Section 4.3 below and Appendix E for details). Beneath each plot are the point estimates and standard errors for select percentiles of these distributions by treatment status. Computed for 2011 to focus on cross-sectional variation.

Panel (b) of Figure 5 plots a similar set of histograms for a different variable of interest,

the expected prescribing share of the focal drug with no meals, $E[s_{aj}(m_{aj} = 0)]$. This helps to solve the puzzle of why the summary statistics showed no clear difference between prescribing patterns for physician-drug 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 drug. Thus, 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 is consistent with a story of meal payments (and the interactions surrounding them) providing information or reminders that counteract potential under-prescribing for some physicians.

Panel (c) of Figure 5 plots the distribution of expected profits 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 profits to drug firms by roughly \$1,624 for the median treated physician. However, our estimates imply that counterfactually extending meals to all untreated cardiologists would have led to net losses, with the incremental profit associated with the median untreated physician estimated to be -\$442. These estimates provide insight regarding why some physicians are targeted by firms and others are not.²⁷

4.3 Supply Model Estimation and Results

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 drug j is removed from the choice set, but prices of the remaining drugs 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 (a^{cs}), the decision error (e^{de}), the manufacturer rebates (r_{jt}), and the marginal costs (mc_j).

To estimate the model for a given vector (e^{de} , T_{jt} , mc_{jt}), we parameterize bargaining ability parameters as a function of drug and regional fixed effects, and specify the econometric

²⁷Liu 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.

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.²⁸ In Section 5, we show welfare for a large range of potential values of λ^{de} . 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 details). We also assume that marginal costs for all jt are equal to 17 percent of the average POS price of generic statins: $mc = 0.17 * pgm$. The value of 17 percent is taken from the average production costs of generic drugs in Sood et al. (2017), assuming that the cost of producing a statin is invariant across molecules and branded/generic status. Appendix H tests robustness to a range of reasonable alternative assumptions regarding (r_{jt}, mc_{jt}) and our results are qualitatively unchanged.

Table 4 summarizes our 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 larger than the weight they place on net costs: $\alpha^{cs} = 1.67$. This may reflect that enrollees are sensitive at the plan choice stage to formulary inclusion of important drugs such as statins. Indeed, Olssen and Demirer (2019) document substantial plan switching based on which statin brands are on formulary in Medicare Part D. It may also capture the role of Medicare Part D program subsidies that limit insurers' financial gains and losses.²⁹

²⁸The 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.

²⁹We do not model such subsidies—e.g., risk corridors and reinsurance—because they are applied at to

Table 4: Supply Parameter Estimates

	Atorvastatin	Lipitor	Lovastatin	Pravastatin	Crestor	Simvastatin
B 2011		0.667 (0.030)	0.060 (0.004)	0.055 (0.004)	0.660 (0.029)	0.045 (0.003)
B2012	0.154 (0.010)	0.755 (0.020)	0.044 (0.003)	0.043 (0.003)	0.771 (0.033)	0.034 (0.002)
^{acs}	1.674 (0.117)					

Notes: $N = 124, 876$ doctor-drug-brand-year observations with standard errors clustered at the doctor level ($Nd = 15, 063$) via delete-120 jackknife and state level via delete-7 jackknife.

5 Equilibrium Welfare Effects of Meals

The above results demonstrate that meals have heterogeneous effects on prescribing, and that they are targeted to more positively 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 our demand and supply parameter estimates to investigate the impact of a counterfactual meal ban on prices, quantities, and welfare in the presence of oligopoly competition, drug firm-insurer negotiated prices, and a range of assumptions regarding potential decision errors in prescribing.

5.1 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 meal 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. We compare the "Ban" scenario to the observed data to understand the full effects of a meal ban—this comparison features prominently in the next subsection on welfare analysis. Our third and fourth scenarios set *out-of-pocket* prices equal to marginal costs with and without a ban ("Ban, p = mc" and "No Ban, p = mc", respectively), allowing us to explore the effects of a meal ban in the absence of a price distortion. These scenarios provide approximations of an "efficient" benchmark—a payment ban and $p^{\text{opp}} = mc$ is efficient at one extreme where $\mathcal{E}^{\text{de}} > 0$, and no ban and $p^{\text{opp}} = mc$ is efficient if \mathcal{E}^{de} is negative and large enough. Table 5 displays several key results from these counterfactuals for 2011. 2012 results are qualitatively similar and shown in Appendix Table A11.

insurers' overall enrollee population rather than at the drug or drug class level.

Table 5: Equilibrium Quantity and Price Effects of Meal Payments (2011)

	Observed	Ban, fix p	Ban	Ban, $p = mc$	No Ban, $p = mc$
Q <i>statins</i>	0.180 (0.001)	0.171 (0.002)	0.172 (0.002)	0.185 (0.002)	0.198 (0.001)
Q <i>Lipitor</i>	0.039 (0.001)	0.032 (0.001)	0.032 (0.001)	0.045 (0.002)	0.054 (0.001)
q <i>Crestor</i>	0.025 (0.000)	0.017 (0.001)	0.017 (0.001)	0.023 (0.002)	0.035 (0.000)
oos <i>statins</i>	18.81 (0.23)	18.81 (0.23)	18.53 (0.23)	2.15 (0.1)	2.15 (0.1)
POS <i>statins</i>	77.40 (0.72)	77.40 (0.72)	75.79 (0.71)	82.41 (1.29)	84.68 (1.40)

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. "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. $N = 124, 876$ doctor-drug-brand-year observations with standard errors clustered at the doctor level ($N_d = 15, 063$) via delete-120 jackknife and state level via delete-7 jackknife.

Table 5 shows total quantities of statins prescribed, OOP prices faced by consumers, and POS prices paid by insurers plus consumers, in the observed and counterfactual scenarios. The price and quantity results in the first three columns highlight several of the issues motivated in Section 3. A ban on meal payments reduces total statin usage as a share of cardiovascular claims by about 0.8 percentage points. This is a four percent reduction in total statin usage. For the focal branded statins, the decrease is 23 percent on average. This is smaller than the large meal treatment effects documented in Section 4.2.3 above because not all doctors receive meals, and some also receive meals from both firms in which case the business stealing effect mutes the overall effect of a meal ban on branded drug usage in equilibrium. The effect of a ban on utilization is similar whether or not we allow prices to adjust because, 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. A ban on meals results in only a small decrease in POS and, in turn, OOP prices.

Turning to the remaining columns of the Table, the quantity estimates also show that pricing above marginal cost reduces total statin usage by about 1.8 percentage points with meals (compare "Observed" to "No Ban, $p = mc$ ") and 1.3 percentage points without meals (compare "Ban" to "Ban, $p = mc$ "). Intuitively, for both Lipitor and Crestor, meals counteract the fact that 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 un-

dershoot the efficient allocation; in the Crestor case, meals cause utilization to fall between the allocations with $p = mc$ with and without meals. The last two columns also illustrate how physician/patient sensitivity to OOP price factors into suppliers' market power—if we counterfactually set $p = mc$ and divorce out-of-pocket prices from point-of-sale prices, point-of-sale prices would increase substantially.

In sum, for the statin market in 2011-12, meal payments from manufacturers to physicians increased demand for branded statins, and thus played an important role in generating profits for the manufacturers involved. They improved allocative efficiency by offsetting the distortion of high branded drug prices, but this was costly to consumers and insurers because promoting branded drugs is an expensive way to increase overall statin usage.

5.2 Welfare Implications of a Counterfactual Meal Ban

To evaluate policies that seek to ban or limit meals and associated interactions, we must quantify how price and quantity effects translate into welfare: consumer, producer, and total surplus. Consumer surplus depends on the extent to which payment effects correct for decision errors that would otherwise lead to underutilization. Motivated by the American College of Cardiology's position that statins are underutilized overall (American College of Cardiology 2017), we suppose in our baseline model that all statins and all physicians are equally subject to a unidimensional decision error \mathcal{L}^{de} that dictates the extent of under- or over-prescribing of statins relative to the outside option. (We explore alternative specifications in Appendix H.) Our counterfactuals allow for more or less substitution to the outside good, but implicitly hold the prices and qualities of the alternative treatments embodied in the outside good fixed.

5.2.1 Welfare Effects as a Function of \mathcal{L}^{de}

In our welfare simulations, we present two different measures of consumer surplus: $CS_{dt}(J)$ accounts for surplus net of out of pocket prices. $CS_{dt}(J_t) = \sum_{ajit} (PpOS(1 - T_{jt}) - Pd_{jt})$ (termed "Consumer Surplus net of transfers" below) further subtracts the portion of drug costs paid by the insurer, which would be consistent with these 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 marginal profit as defined in Section 3.3, minus salesforce and meal costs: $PS_{jt} = \mathcal{L}^{de} \wedge (PpOt(1 - T_{jt}) - mc_{jt}) - C_j \wedge^{-1}$, where T_{jt} is the manufacturer rebate, mc_{jt} captures the cost of manufacturing and distributing the marginal

unit of drug j , and $C^{\text{TMd}j=1}$ is the total cost of a meal-based relationship with a physician.³⁰ Total Surplus is the sum of Producer Surplus and Consumer Surplus net of transfers. 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.

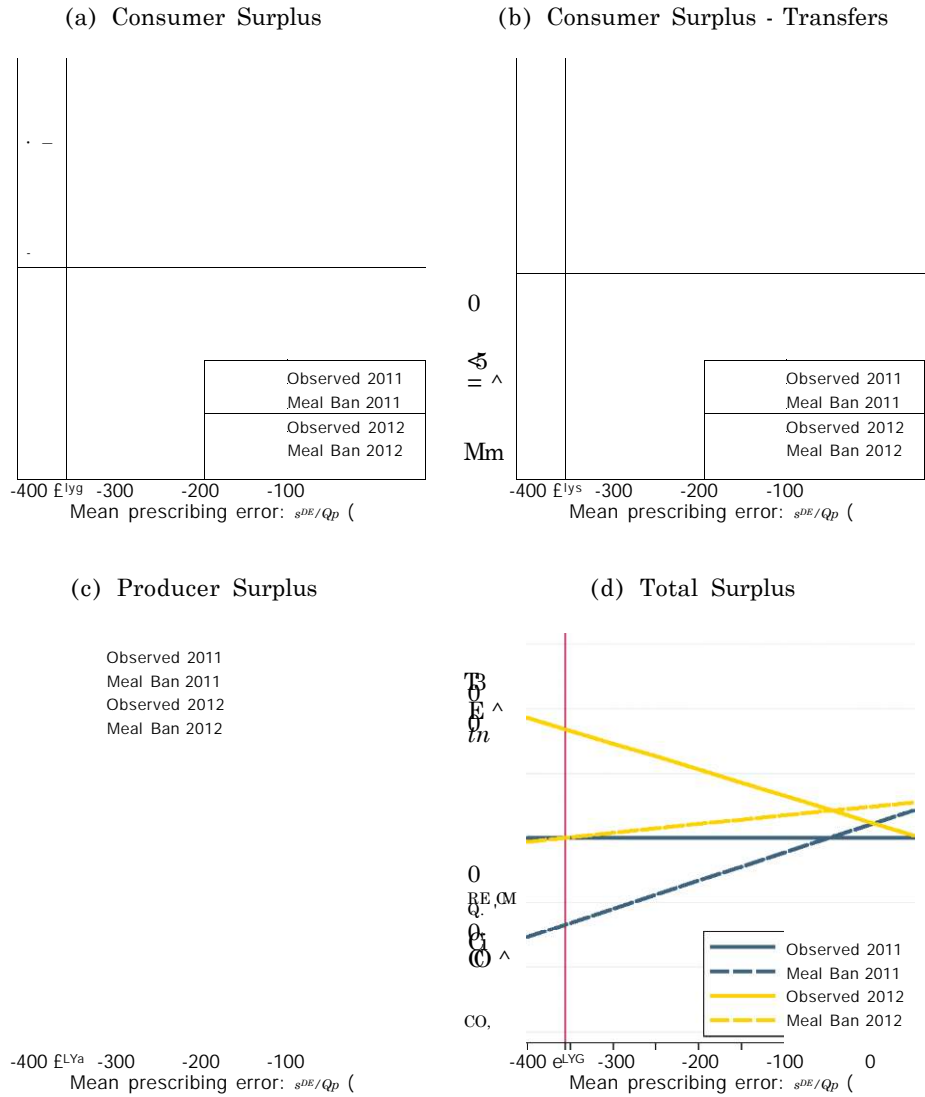
The results are summarized graphically in Figure 6 with all measures represented as percent changes relative to the baseline of the observed outcomes in 2011. Thus, "Observed 2011" is a flat line at zero, we compare "Observed 2012" to "Observed 2011" to quantify the welfare effect of atorvastatin entry, 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) is decreasing in \mathcal{L}^{de} . Intuitively, the more negative \mathcal{L}^{de} is, the greater the implied benefit of taking statins, and in turn of the statin market expansion in 2012. At $\mathcal{L}^{\text{de}} = 0$, Consumer Surplus increased by \$0.91 per patient due to generic atorvastatin entry, but Consumer Surplus net of transfers increased even more (\$3.17) 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. More negative \mathcal{L}^{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 justify the increased expenditures. Allowing meals improves consumer surplus in 2011 for $\mathcal{L}^{\text{de}}/0^p < -\72 , but only improves consumer surplus net of transfers for the more extreme threshold value of $\mathcal{L}^{\text{de}}/0^p < -\268 . In 2012, with a single branded firm offering meal payments, and another generic option available in atorvastatin, the decision error necessary to justify meals is less extreme, with cutoffs near $-\$68$ for pure Consumer Surplus and $-\$204$ for Consumer Surplus net of transfers. Intuitively,

³⁰As discussed previously, 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 in Liu et al. (2020). See Appendix E for construction of baseline and alternative rebate assumptions. As shown in Appendix H, Producer Surplus changes under alternative cost and rebate assumptions, but Consumer Surplus is largely unchanged.

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. Results shown for $e^{de} \in [-\$400, \$50]$. Detailed results for select values of e^{de} available in Appendix Table A12.

the large negative cutoffs of -\$268 and -\$204 reflect that, with our specification of a flat \mathcal{L}^{de} across all drugs, the market expansion effect of meals on overall statin use (which is 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. 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 effect 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 $\mathcal{L}^{de}/g^p < -\49 . If there is no underlying decision error, a meal ban increases Total Surplus by \$0.32 per patient. Alternatively, if decision errors are equivalent to the average effect of meals on revealed willingness to pay ($\mathcal{L}^{de}/0^p = -\41), such that meals cancel out decision errors on average among those receiving them, 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, the extent to which insurers internalize cardiologist decision errors ("Pricing"), and decision errors correlated with meal responsiveness $\mathcal{L}^d = Y^{de} *$ instead of fixed 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 \mathcal{L}^{de} . The effects of a meal ban are qualitatively and quantitatively similar across all modeling assumptions, with the exception that $\mathcal{L}^d = Y^{de} *$ implies that meals are more likely to be welfare improving (e.g., Total Surplus increases for $\mathcal{L}^{de} < -\$25$, vs. the lower threshold of -\$49 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 pattern is robust: the effect of meal payments on consumer welfare depends critically on the extent to which meals simply increase usage vs. increase usage in cases where the drug would be severely underutilized $\mathcal{L}^{de} \ll 0$.

5.2.2 Calibrating the Decision Error Magnitude using Clinical Data

The extent of over- and under-utilization (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 λ^{de} is sufficiently negative 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 \$516.³¹ This measure suggests that the $\lambda^{de}/9^p$ 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 -\$268 at which Consumer Surplus net of transfers improves in the presence of meals.

This calibration exercise is instructive if 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 75 percent of (randomly selected) Medicare patients in our sample would experience the clinical benefits of statins from the medical literature, even if the remaining 25 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

³¹To 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.

estimates for a given sample. For illustrative purposes, we have done so at the aggregate 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 and patient level clinical trial 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 may also 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 and in part due to the difficulty of mapping any estimated effects into welfare in light of other market frictions. These issues are particularly salient in light of recent policy debates over conflicts of interest in the U.S. health care and financial services industries.

We propose a framework to address these challenges and implement it 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 for statins in 2011-12, 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, and that firms target payments to physicians who will be responsive to their interactions and do not target those who do not appear to be 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 under-prescribing, 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. When we calibrate the decision error parameter to clinical data on the value of statins, we find that, in our estimated model for statins, meal payments increase consumer surplus as well due to under-prescribing at baseline. The magnitudes of these effects are large in the sense that they are similar to the introduction of generic atorvastatin, one of the largest generic introductions of all time.

There are limitations to our approach. We focus our case study on a particular market, cardiologists and statin prescriptions in the Medicare Part D program, 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. The framework we have developed here could be a starting point for these explorations.

A study that measures responses of prescribing to payments across many drugs could be especially useful. Such a paper will likely require alternative strategies for estimating price elasticities in other contexts, and adaptations to allow for new drugs or other cases where the information environment might be changing during the time frame studied. Finally, scaling our decision error calibration approach for mapping effects into welfare would require careful analysis of the suitability of available clinical evidence.

Can our current set of 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 more educational interactions and information exchange between industry and physicians.

Much can be gained from future research looking at similar phenomena in different contexts. In our results here, 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 also be used in many other 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.

Finally, we 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., 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.
- Ahrens, A., Hansen, C. B., and Schaffer, M. (2020). `pdslasso` and `ivlasso`: Programs for post-selection and post-regularization ols or iv estimation and inference.
- 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., Chen, D., Chernozhukov, V., and Hansen, C. (2012). Sparse models and methods for optimal instruments with an application to eminent domain. *Econometrica*, 80(6):2369-2429.
- 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 Waldfoegel, 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.
- Besanko, D., Dranove, D., and Garthwaite, C. (2020). Insurance access and demand response: Price and welfare implications. *Journal of Health Economics*, 73.
- Bhattacharya, V., Illanes, G., and Padi, M. (2020). Fiduciary duty and the market for financial advice. Working paper.
- 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., Dube, 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.
- CECS (2012). The dartmouth atlas of health care.
- 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.
- CMS (2012). Centers for medicare and medicaid services: Research, statistics, data and systems.
- CMS (2013). Centers for medicare and medicaid services: Provider data.
- 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).
- Collard-Wexler, A., Gowrisankaran, G., and Lee, R. S. (2017). Nash-In-Nash bargaining: A microfoundation for applied work. *forthcoming in Journal of Political Economy*.
- ConsumerReports (2014). Evaluating statin drugs to treat: High cholesterol and heart disease.
- Crawford, G. and Yurukoglu, A. (2012). The welfare effects of bundling in multichannel television. *American Economic Review*, 102(2).
- 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. (2017). Physician vs. patient incentives in prescription drug choice. Working paper.
- 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.
- Harvard Men's Health Watch (2014). Help for your cholesterol when the statins won't do. https://www.health.harvard.edu/newsletter_article/help-for-your-cholesterol-when-the-statins-wont-do <accessed 01/15/2021>.

- 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 Lee, R. S. (2017). Insurer competition in health care markets. *Econometrica*, 85(2):379-417.
- 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.
- 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.
- 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.
- Olsen, 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.
- 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?

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 may in general depend on 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 regulates 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.³²

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. 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.³³

A.2 Regional Prices and Formulary Variation

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

³²In 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").

³³Partial 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.

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.³⁴

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 national drug code (NDC). 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) 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.³⁵ 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.³⁶

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, which 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

³⁴In 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.

³⁵<https://q1medicare.com/PartD-The-2014-Medicare-Part-D-Outlook.php>

³⁶ <https://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/MCRAdvPartDEnrolData>

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.

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. For generic atorvastatin in 2012, there was significant variation in terms of both point-of-sale (CV =0.11) and out-of-pocket price (CV =0.24). This price variation, at the state-year-drug level, is presented for our focal drugs in Table A1 below.

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

		2011				2012				2011 -2012	
		mean	s.d.	[^] cross-sec.	s.e.cross	mean	s.d.	[^] cross	s.e.cross	[^] panel	s e 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- -statin	OOP					10.03	2.420	0.850	0.300		
	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").

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 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 (*ft* 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 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-drug-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-drug level (dj ; Panel b) only for Crestor and Lipitor and uses the sample corresponding to column (4).

B.2 Linking Payments Data

The payment data is based on publicly available data released by firms prior to the Sunshine Act, which began requiring reporting in 2013. In the data, physician-level identifiers were often limited to a name, city of address, and perhaps a specialty. Back when the reports were still posted on firms' websites, the enterprise software company Kyruus collected them as a part of their initiative to analyze physician-firm relationships. 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.

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

Table A2: Sample Descriptions

		(1)	(2)	(3)	(4)
Panel (a): Unit-of-analysis (<i>djt</i> , 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 <i>djt</i> obs.	unique	217,987	158,939	123,809	121,121
Panel (b): Cardiologist-drug (<i>dj</i> , 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
	Any type	0.293	0.326	0.350	0.352
Pfizer — Lipitor	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
	unique	39,634	28,898	25,323	25,156
N <i>dj</i> obs.	unique	39,634	28,898	25,323	25,156
Panel (c): Cardiologist (<i>d</i> , Crestor and/or Lipitor)— Payment data					
Cardiologists (N <i>d</i> 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) includes all cardiologists in Physician Compare; (2) restricts sample to those in (1) with >500 cardiovascular claims in both 2011 and 2012; (3) restricts sample to those in (2) with non-zero *qdjt*; (4) restricts sample to subset of cardiologists in (3) for which we estimate \hat{d}_j , which are then used in the MTE estimation.


payments. Within general payments we identify three sub-categories: "meals," "travel or lodging," and "consulting, speaking or education."

C Additional Theory and Connection to Empirics

C.1 Graphical Framework

To build intuition regarding the potentially complex effects of payments in the presence of other frictions, consider a simple model where payments shift the demand curve outward. Panel (a) of Figure A1 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 $Q^{eff}Q^b$ —marginal patients prescribed the drug in the presence of payments to physicians receive negative health benefits.

Figure A1: Welfare Analysis with Other Frictions



In a setting with perfect competition, this conceptual framework suggests that the causal effects of payments on prescribing are 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 drugs 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 of this model is presented in Panel (b) of Figure A1: 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 $Q^m < Q^e < Q^b$ (pictured) or cause prescribers to overshoot the optimum $Q^m < Q^e < Q^b$. 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 A1. 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: $Q^m < Q^{ins} < Q^b$.³⁷ 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. 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 health care 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 payments); 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 A1 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 $Q^{de} < Q^b < Q^{eff}$. The next Section describes how our welfare analysis incorporates a "decision error" parameter that allows for a range of assumptions on

³⁷Another 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.

how payments might counteract, overshoot, or reinforce any baseline biases.

C.2 Consumer Surplus

As outlined in Section 3, 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 allow for the demand unobservable to have two components:

$$d_{jt} = d_{jt} + e_{jt}$$

where d_{jt} is a typical demand unobservable that impacts both choices and true realized utility, but e_{jt} is a "decision error" that impacts decision utility but not realized utility.

Given this model, we represent expected consumer surplus as:

$$CS_{dt}(t) = \sum_j Q_{dt} \ln \left(\frac{v_j + (E_{jt})^{1-\alpha}}{V_j} \right) - \sum_j E_{jt} \quad (1)$$

CS implied by decision utility adjustment for "decision errors" and meals

The second term adjusts consumer surplus for the presence of a decision error that results in under- ($e_{jt} < 0$) or over-prescribing ($e_{jt} > 0$) as in Figure A1(d) above, as well as the countervailing (or reinforcing) effect of meals.

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, $e_{jt} = 0$ would be a case with no decision error at all, where meals simply bias utilization of promoted drugs upward. By contrast, $e_{jt} = -\delta_j$ would be a case where meals perfectly correct prescribing errors among those who receive them.

We study the welfare implications of two different specifications of decision errors. In our main specification, we set a constant decision error across all doctors and statins $e_{jt} := e^{de}$, and we simulate counterfactuals for a range of decision errors, from substantial underprescribing to overprescribing.³⁸ 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*

³⁸This 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.

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 $e^d := Y^{de} Q_j$, varying scalar Y^{de} to again allow for a range of potential under- or over-prescribing in the absence of meals. In this specification, $Y^{de} = -1$ 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 $Y^{de} = 0$, undercorrect $Y^{de} < -1$ or overcorrect $Y^{de} > 0$ 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.

C.3 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, drug 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:

$$j - mc_j (E[q_j = 1 | I_{dj}]) > R_j \quad (10)$$

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 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 $C_j^{mdj=1}$ 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 $E[q_j^{-1} - q_j^{\text{no}} | I_j]$ will be a function of the expectation of total size of the doctor's patient flow Q_d and the choice probability function as given in (3). In particular, it will be a function of the expectation of the parameter θ^m which determines the effect of the meal interaction on the mean utility weight the doctor assigns to drug j .

C.3.1 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 drug j whenever

$$j - m c_j \left[q_m^{\text{no}} - q_m^{\text{meal}} \right] > c_m^{\text{no}}(N_{jrd}, 0) - c_m^{\text{meal}}(j, 0). \quad (11)$$

To deconstruct this expression, we use dq/dN_{jrd} as an approximation to $(q^{\text{meal}} - q^{\text{no}})$.³⁹

We also specify a particular cost function $C_{dj}(N_{jrd}, 0) = \theta A_{dj} N_{jrd}^\alpha$. Here A_{dj} represents an access cost shifter that may be drug-doctor specific, N_{jrd} 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 $\alpha > 1$. Here we also use dC/dN as an approximation to $c_m^{\text{meal}}(N_{jrd}, 0) - c_m^{\text{no}}(N_{jrd}, 0)$.

Substituting these values gives

$$j - m c_j Q_{dj} > A_{dj} N_{jrd}^\alpha. \quad (12)$$

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

$$\ln(Q_{dj}) + \frac{1}{\alpha} \ln(A_{dj}) - \frac{1}{\alpha} \ln(N_{jrd}) + \ln(j - m c_j) + \ln \left(\frac{dQ_{dj}}{dN_{jrd}} \right) > 0. \quad (13)$$

$f(X_j; \hat{\theta})$ $g(Z_{dj}; \hat{\theta})$ $\hat{\epsilon}_{dj}$
 flexible approx. via Lasso residual: correlated with $d^{\text{TM}} + \epsilon_{dj}$

³⁹For our primary demand specification, this partial derivative is given by: $Q_{dj} \frac{dQ_{dj}}{dN_{jrd}} = \left(\frac{dQ_{dj}}{dN_{jrd}} + \frac{dQ_{dj}}{dN_{jrd}} \right) (1 - \alpha)$

C.4 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:

$$p_{jt} = \arg\max_{p_{jt}} (n_j \cdot s_j \cdot m_{jdt})^{b_j} (CS_t(J_t) - CS_t(J_t \setminus j))^{1-b_j}$$

$$m_{jdt} + b_{jt} \frac{dq_{jt}}{q_{jt}} = \frac{dq_{jt}}{q_{jt}} \frac{CS_t(J_t) - CS_t(J_t \setminus j)}{CS_t(J_t) - CS_t(J_t \setminus j)} + p_{jt} \theta_j (1 - m_{jdt}) / (1 - T_{jt})$$

where $q_{jt} := \sum_d q_{jdt}$ denotes the sum over physicians. The term b_{jt} 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 drug j : $CS_t(J_t) - CS_t(J_t \setminus j)$. T_{jt} reflects the rebate rate to insurers off the posted price p_{jt} .

C.5 Alternative Model of Insurer Pricing

In our main specification, we assume that insurers negotiate drug prices as a function of consumer surplus as implied by decision utility. In reality, it is possible that insurers are aware of over- or underprescribing of some drugs and take that into account when negotiating prices. In such a scenario, we would want to replace CS_{dt} in equation (6) with the following:

$$CS_{dt}(J_{dt}) := a^{de} \left[\sum_j Q_{jdt} (1 + (\lambda_j e^{-1-A_j})^{-a^{de}}) \right] + \sum_j \left(\frac{Q_{jdt}}{Q_{jdt}} \right) \cdot \left(\frac{Q_{jdt}}{Q_{jdt}} \right)$$

CS implied by decision utility adjustment for "decision errors" and meals

$$q_{jdt} (1 - T_{jt}) \cdot p_{jt} D_j$$

insurer drug costs

Here, we include a parameter $a^{de} \in [0,1]$ that allows for a range of assumptions regarding how insurers incorporate decision errors and meals into their surplus measure. This model accommodates the fully "naive" case where insurers negotiate prices under the assumption doctors know best ($a^{de} = 0$), the fully "sophisticated" case where insurers perfectly adjust consumer surplus for decision errors and meals ($a^{de} = 1$), and every case in between.

We present results for the fully "naive" case in the main text, and for the fully "sophisticated" case in Appendix Tables A13 and A14. As shown in Appendix Table A13, if insurers

are sophisticated ($\alpha^{de} = 1$), then the effect of a meal ban on prices depends on the decision error. A ban leads to Crestor prices decreasing by \$2 if $e^{de} = -350$ and *increasing* by \$2 if $e^{de} = 0$. 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. As shown in Appendix Table A14, the value of α^{de} has little impact on the welfare implications of a meal ban.

D Parameter Estimation Routine

The following outline details the steps necessary to recover the demand parameters (θ^p , X , β_j , 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). Note that for ease of notation, here we use θ^m to refer to the distribution of meal treatment effects.

Jackknife doctors and create bootstrap samples (main source of variation for inference, blocked to cluster at doctor level)

1. Replicate the full sample of djt-level observations 500 times, dropping the observations for a randomly selected VWd doctors; samples indexed by k

For each k , perturb quantities (allows for sampling error in prescribing shares in first step of demand estimation)

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

For each k , estimate price, nest and V_j parameters

5. Estimate Eq. (8) to recover price (Q^{p^k}) and nest (A^k) parameters, and doctor-molecule fixed effects (v_j)
6. Parameter estimates: for θ^p and A , point estimates are from estimation on the full sample, standard errors are given by the standard deviation across the 500 k samples
7. Shrink each V_j estimate towards the j -specific mean using the standard deviation of V_j across the 500 k samples as the standard error in the standard empirical Bayes shrinkage formula

For each k , estimate meal parameters

8. Keep V_j 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, d^{m^k} as described below in Section D.2 and based on Eqn. (9)
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 Appendix 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) \hat{y} on the vector of possible controls (and transformations thereof) X , with the selected variables given by $L(X)$
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 $Q^{m>k>s}$ (and other MTE parameters) using the variables selected in the opposite sub-sample s' taking the union of $L(X)$ and $L(X)^{m's}$ for controls
5. Solve for the k -specific estimate $d^{m^k} = (d^{m^kA} + 0^{m^kB})/2$
6. Parameter estimate: for 6^m (and all corresponding MTE-based estimates) point estimates and standard errors are given by the median, $Q^m = \text{median}(9^{m^k})$, and the "median deviation", $\text{s.e.}(0^m) = \sqrt{\text{median}(9^{m^k} - 9^m)}$, 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-specific responses to the instruments, all Z variables are interacted with drug-specific dummies.⁴⁰

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 $\text{MaxPenaltyGuess} = 2 \times \max(x'y)$, where x is the pre-standardized regressor matrix and y is the vector of the outcome variable, from a minimum of $[\text{MaxPenaltyGuess}/1000]$; 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.

⁴⁰We obtain qualitatively similar results if only using the levels of X variables as controls and/or not allowing instruments to be drug-specific.

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 denoting unobservables in the outcome and treatment equations, respectively:

$$\begin{aligned}
 Yd &= f^d(Xi) + Wd && \text{for } d = 0,1 \\
 Y &= DY^1 + (1 - D)Y^0 \\
 D_i &= 1\{g(X_i, Z) > V_i\}.
 \end{aligned}
 \tag{14}$$

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

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

where U , the unobserved resistance to treatment, is given by the quantiles of V .⁴¹

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 linear probability model 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 $E[W^1 - W^0 \mid U_i = u]$).

As shown by Andresen (2018), posterior estimates of doctor-specific treatment effects can

⁴¹This smoothing creates the unit interval that is the x-axis for all MTE curves.

be calculated using the following formula:

$$\begin{aligned}
 E[Y^1 - Y^0 \mid X = x, Di = d, Pi = p] &= x(P^1 - P^0) \\
 &+ dE[W^1 - W^0 \mid Ui < p] \\
 &+ (1 - d)E[W^1 - W^0 \mid Ui > p],
 \end{aligned}
 \tag{16}$$

where P_i is the doctor's propensity score.

D.3 Important Variables & the Importance of Many Variables

Table A3 reports the top variables selected by the Lasso—specifically, the number of subsamples within which the variable is selected in either the outcome or treatment 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

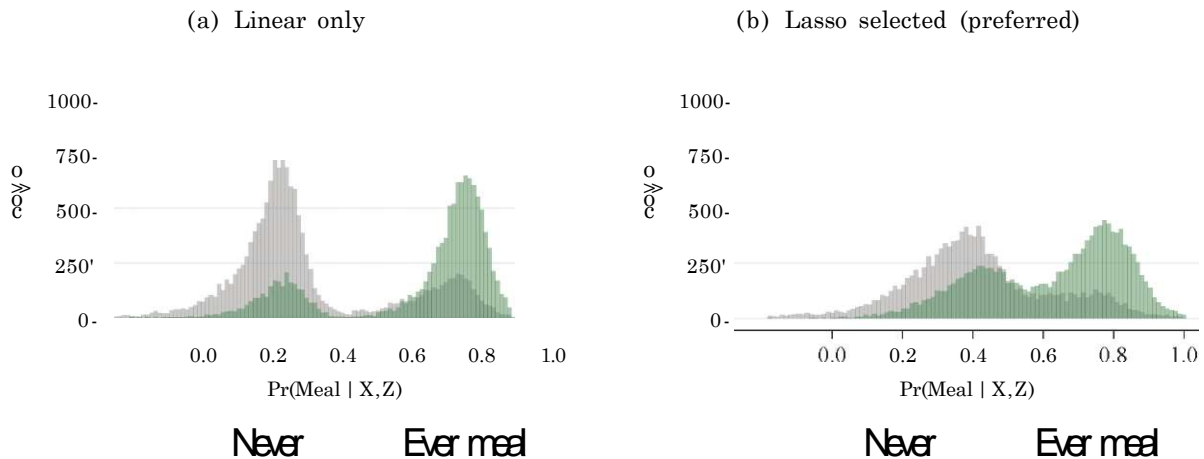
	Num. subsamples selected in
Controls (X)	
log(Doc. annual avg. cardiov. claims)	1,000
log(Hosp. num. admissions)	996
log(HSA % uninsured)	993
Hosp. annual doc.-avg. cardiov. claims	981
log(HSA annual doc.-avg. cardiov. claims)	980
log(HSA Medicare Advant. eligibility)	980
Ad spend ²	978
log(HSA % Medicaid)	975
HRR Share AAMC faculty	975
log(HRR % uninsured)	964
log(Nearest AMC drive time)	963
State low-income subsidy enroll.	957
Doc. graduation year	955
HRR num. AAMC faculty	949
log(HSA share AAMC faculty)	947
AMC AMSA Instruments (Z)	
Lipitor, HRR AMSA	1,000
Lipitor, HSA AMSA	955
Crestor, HRR AMSA X faculty-wgt.	886
Lipitor, HRR AMSA X faculty-wgt.	819
Crestor, HSA AMSA X faculty-wgt.	658
Crestor, HRR AMSA	645
Lipitor, HSA AMSA X faculty-wgt.	603
Crestor, HSA X faculty-wgt. X drive time	426
Crestor, HRR X drive time	391
Lipitor, HSA X drive time	384

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 A2 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.

Figure A2: 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).

D.4 Perturbation and Shrinkage

We are concerned that the doctor-molecule mean utility parameters (the fixed effects) might be influenced by noise (since we only observe two years of utilization), especially for 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 then resample at the use case 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.

Ultimately, about 94 percent of the observations are shrunk by less than 10 percent of their raw values, and whether or not we perturb market shares in this way does not substantially alter our demand estimates (see Table A9 below).

E Role of Rebates

The negotiation modeled in Section 3.3 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). Like nearly all pharmaceutical research, 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 j_l that suppliers receive, and that insurers pay after cost-sharing is applied, are net of rebates T . 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 5 plots the distribution of expected changes in firm revenue from targeting meal payments. 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 T and "other suppliers' markup" w , so that manufacturer revenue becomes $R(j_l) = \sum_{j,l} q_{j,l} p_{j,l} (1 - w - T)$. In the supply side estimation, welfare simulations, and simulations of manufacturer revenue, we rely on researchers' estimates of T and w , and we test the sensitivity of our results to our decisions on how to use these estimates.

To obtain the components T and w , we rely on multiple sources.⁴² CMS has reported

⁴²We are aware of several sources of information on $T+w$: 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 $p^{mfr} = p^{pos} * (1 - w - T) = 0.673 * p^{pos}$, suggesting $T + w = 0.327$. Kakani et al. (2020)'s estimates suggest

average manufacturer rebate percentages overall ($T = 0.175$) and for cardiovascular drugs specifically ($T = 0.263$) going back to 2014.⁴³ Arcidiacono et al. (2013) assume $T = 0.151$ 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 $T + w = 0.32$ 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 $T + w = 0.263 + 0.190 = 0.453$ would be more appropriate. For our supply side estimation and counterfactuals, the above papers suggest $T = 0.32 - 0.037 - 0.152 = 0.131$ 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 $T = 0.263$ based on cardiovascular drugs only. We use ($T = 0.263, T + w = 0.453$) in the main text. As a robustness check, we consider ($T = 0.131, T + w = 0.32$) 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 ($T = 0.483, T + w = 0.673$) (based on Aitken et al. (2018); Arcidiacono et al. (2013)); we stick with the alternative assumptions ($T = 0.131, T + w = 0.32$) in our robustness check below. Finally, for generic drugs, we rely on Sood et al. (2017), which is the only source explicitly breaking out generics, and assume ($T = 0.24, T + w = 0.41$).

Comparing Table A4 to the results from our preferred specification in the main text (Figure 4, Panel c), this alternative (larger) pass-through assumption 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.

an average $T + w = 0.32$ across a wide range of drugs that (unfortunately) explicitly excluded statins. Sood et al. (2017) suggests $T + w = 0.42$ across a range of branded and generic drugs.

⁴³https://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	-20.9 (794.0)	546.6 (869.6)	1,573.8 (1335.3)	3,525.8 (2318.7)	6,544.8 (3647.2)
Ever	1,203.3	2,153.9	4,216.7	7,960.9	13,320.4
meal	(239.7)	(389.4)	(705.0)	(1,231.1)	(1,944.3)

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 are 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. 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 indicated patients 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, for a patient who is clinically indicated, is $\$19,147.50/37.13 = \516 .

Mapping these numbers into our demand analysis requires taking a stand on the extent to which the patients being prescribed statins are indeed clinically indicated. 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 start from the assumption that patients being prescribed statins in the Medicare population in our analysis are indeed clinically appropriate. However, we also discuss how our results can be recalibrated using X percent of patients as appropriate, for the reader's preferred $X \in [0,100]$.

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 and instrument sets in the Lasso regressions. Both tables are based only on doctor-drug observations for 2011 and either Lipitor (Pfizer) or Crestor (AstraZeneca).

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

	mean	s.d.	f meal	f cardiov. share
<u>Doctor level:</u>				
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
Num. specialties	1.40	0.550	-0.031	-0.0042
Female = 1	0.088	0.284	-0.096	0.016
Drive time to nearest AMC (sec.)	877	1,900	-0.016	-0.053
AAMC Faculty = 1	0.096	0.295	-0.226	0.120
AMSA faculty wgt.	2.27	7.46	-0.234	0.118
<u>Hospital level:</u>				
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. AMCs	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
<u>Zipcode-drug level:</u>				
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: participates in Medicare's electronic prescribing incentive program; PQRS: participates in Medicare's Physician Quality Reporting System; EHR: participates in Medicare's electronic health records incentive program. The "f meal" and "f cardiov. share" columns report the coefficient from a regression of either a dummy for meal receipt (f meal) or the standardized cardiovascular share of a drug (f cardiov. share) regressed on the standardized variable.

Table A6: Summary of Potential Controls and I.V.s—HSA, HRR, and State

	mean	s.d.	f meal	f cardiov. share
<u>HSA level:</u>				
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. AMCs	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
AMSA avg.	25.8	2.07	-0.094	0.040
AMSA faculty wgt.	3.59	1.93	-0.102	0.142
<u>HRR level:</u>				
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. AMCs	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
AMSA avg.	25.6	3.16	-0.034	-0.0098
AMSA faculty wgt.	2.45	1.29	-0.086	0.122
<u>State level:</u>				
Plan enrollment	992,164	689,303	0.070	0.062
Plan enroll., low income subs.	410,867	301,107	0.069	0.078

Notes: HSA-level: summarizes HSA-level aggregates for each physician, excluding physician's affiliated hospital. HRR-level: summarizes HRR-level aggregates for each physician, excluding the HSA of physician's affiliated hospital. The "f meal" and "f cardiov. share" columns report the coefficient from a regression of either a dummy for meal receipt (f meal) or the standardized cardiovascular share of a drug (f cardiov. share) regressed on the standardized variable.

G.2 HSA-level First Stage for Meals

Table A7 replicates Figure 2 Panel (c) 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 IV First Stage

	(1)	(2)	(3)	(4)
AMSA avg.	-0.0158 (0.0011)	-0.0158 (0.0012)	-0.0146 (0.0012)	-0.0145 (0.0012)
AMSA avg. X time to nearest AMC		-0.0002 (0.0005)		-0.0003 (0.0005)
AMSA avg. X faculty-wgt.			-0.0062 (0.0015)	-0.0062 (0.0014)

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 AMC 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 IVs on treatment propensity) and the "reduced form" relationship (direct effect of the IVs 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 of cardiologists 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

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 first stage regressions using meal payments as the dependent variable, and reduced form regressions using cardiovascular shares (s_{djt}) and the doctor-molecule mean utilities defined in Section 4.1.1 (ip_{dj}) as dependent variables.

To conduct these tests, we use the same bootstrap and Lasso-based approach as in all main specifications, but we also interact the selected instruments with an indicator for whether or not the observation belongs to a particular placebo sample. Then, instead of reporting the coefficient on every instrument, we report the average change in the dependent variable that would be associated with a 1 s.d. change in all AMC AMSA scores given the estimated coefficients on the instruments for the placebo and non-placebo 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), we always find that in the non-placebo subsample, the instruments behave as we expect, leading to fewer meals and corresponding declines in both cardiovascular shares and the doctor-molecule mean utilities (Cols. 1, 3, 5). In the placebo subsamples (Cols. 2, 4, 6), we consistently estimate smaller and/or statistically insignificant first stage and reduced form relationships. These tests are not high-powered, due to the relatively small subsamples that are identifying the placebo-specific effects. Still, these results—that, among cardiologists less affected by AMC policy spillovers, such spillovers are not predictive of either meal payments or our dependent variables—are reassuring regarding our exclusion restriction.

G.4 Demand Estimation Results: Details and Robustness

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.

Table A8: First Stage and Reduced Form Estimates—Placebo Tests

	FS: $1_m > 0$		RF: s_{djt}		RF: \hat{d}_j	
	Non-plac. (1)	Plac. (2)	Non-plac. (3)	Plac. (4)	Non-plac. (5)	Plac. (6)
Panel (a): Placebo in States with Payment Restrictions						
SD.V.X100 SAMSA I.V. p-value	-5.796 [<0.001]	2.618 [0.857]	-0.140 [<0.001]	1.072 [0.059]	-1.427 [0.003]	22.195 [0.134]
mean D.V.X 100	52.99	6.29	3.88	4.30	-226.74	-208.27
avg. N obs.	24,201	349	24,201	349	24,001	346
Panel (b): Placebo with High own-AMSA Faculty						
SD.V.X100 SAMSA I.V. p-value	-5.824 [<0.001]	0.810 [0.208]	-0.130 [<0.001]	0.378 [0.337]	-1.240 [0.003]	12.039 [0.502]
mean D.V.X 100	52.96	25.61	3.86	4.87	-226.94	-207.69
avg. N obs.	23,979	570	23,979	570	23,782	565
Panel (c): Placebo in High own-AMSA Hospitals						
SD.V.X100 SAMSA I.V. p-value	-5.559 [0.003]	-1.416 [0.011]	-0.125 [<0.001]	-0.046 [0.006]	-1.123 [0.003]	-1.753 [0.271]
mean D.V.X 100	53.6	25.14	3.83	5.17	-227.55	-203.51
avg. N obs.	23,456	1093	23,456	1093	23,263	1084
Panel (d): Placebo with $E(1_m > 0 X) < 0.1$						
SD.V.X100 SAMSA I.V. p-value	-5.823 [<0.001]	-4.442 [0.21]	-0.128 [<0.001]	-0.043 [0.026]	-1.237 [0.006]	2.652 [0.276]
mean D.V.X 100	53.75	8.32	3.84	5.49	-227.47	-195.84
avg. N obs.	23,781	768	23,781	768	23,586	761

Notes: Reports the first stage (FS) and reduced form (RF) OLS regressions using the same 500 bootstrap variable selection and estimation procedure used in the main specifications. $\Delta_{\text{SAMSA I.V.}}$ indicates the change in the dependent variable (D.V.) if all AMCs increased their Col policies by one standard deviation per the AMSA scores (multiplied times 100 for ease of viewing). The FS D.V. is meal probability. In RF: s_{djt} , the D.V. is the 2011 focal drugs' share of cardiovascular claims, and in RF: \hat{d}_j , the D.V. is the estimated doctor-molecule intercept from the price and nest regression (\hat{d}_j). p-values are based on a test of joint significance of the set of selected instruments. Within each panel, each pair of columns (1-2, 3-4, 5-6) reports the estimates of $\Delta_{\text{SAMSA I.V.}}$ from a single regression where the instruments are interacted with an indicator for the placebo sample indicated in that panel (e.g., "States with Payment Restrictions", "Hi-AMSA Faculty", etc.). The mean D.V. and the average number of observations across the 500 bootstrap samples are also reported.

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)
Q^p	-0.00753 (0.00016)	-0.00708 (0.00013)	-0.00033 (0.00002)	-0.00753 (0.00003)	-0.00757 (0.00017)	-0.01316 (0.0001)	-0.00744 (0.00017)
A_{statin}	0.423 (0.011)	0.279 (0.012)	0.965 (0.001)	0.423 (0.002)	0.426 (0.011)		0.441 (0.011)
$A_{\text{strong statin}}$							-0.026 (0.008)
mean(n^p)	-0.21	-0.16	-0.14	-0.21	-0.21	-0.22	-0.21
s.d.(n^p)	0.23	0.17	0.16	0.23	0.23	0.23	0.23
F-stat.	476.9	655.2	14980.3	476.9	402.2	45652.6	344.6
mean($\hat{d}_j Q^p $) $_{\text{str. statins}}$	-296.7	-350.9	-4634.3	-296.7	-296.8	-211.2	-298.1
mean($\hat{d}_j Q^p $) $_{\text{other gen.}}$	-315.7	-368.2	-4730.2	-315.7	-315.6	-228.6	-315.8
s.d.($V_j Q^p $)	74.6	91.2	1161.8	74.6	73.8	63.1	74.3
$R^2(s_{ajt} \cdot V_j)$	0.287	0.304	0	0.287	0.290	0.293	0.289
$R^2(s_{ajt} \cdot V_j + Q^{pp})$	0.425	0.408	0.011	0.425	0.430	0.425	0.426
$R^2(S_{ajt} \cdot V_{dj} + Q^{pp})$	0.809	0.834	0.147	0.809	0.807	0.881	0.803
Specification							
Instrument Q^p	Y		Y	Y	Y	Y	Y
Instrument A	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. 8. Standard errors for the main parameters (Q^p and A), in parentheses, are based on the standard deviation of the 500 point estimates from the perturbed-bootstrap samples. $R^2(s_{ajt} \cdot \cdot)$ reports the R^2 from a regression of $s_{ajt} = \ln(s_{ajt}/s_{dot})$ on some combination of the molecule- (V_j) or molecule-doctor-level (V_{dj}) fixed effects, and possibly the price effect (Q^{pp}).

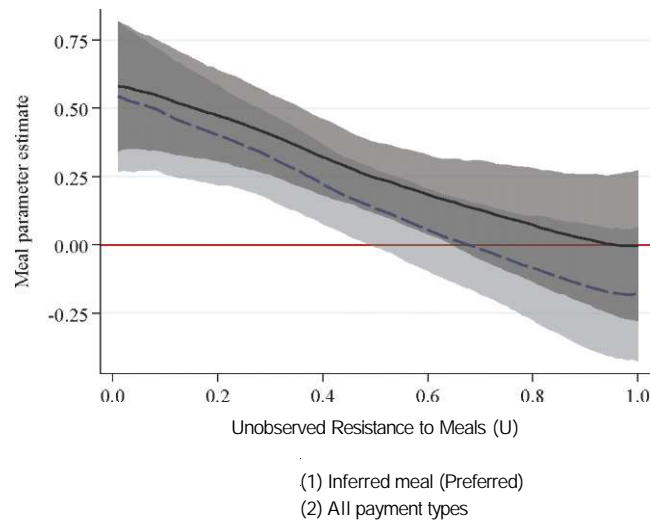
G.5 Exploration of Treatment Effects of Payments

In the following analyses, we explore alternative indicator variables for payments, an alternative approach to capturing business stealing, the potential existence of intensive margin (dollar value) meal effects, and the potential existence of within-practice spillover effects of meals.

Figure A3 plots MTE estimates of 9^m and, in the table below, the ATE and LATE implied by the MTEs, using: our preferred indicator for meal-based relationships (Col. 1, "Preferred"); and an indicator for receiving any kind of payment (e.g., meals, consulting, speaking, travel, or research) from the firm (Col. 2). The results are somewhat smaller for "all payment types," but are not statistically significantly different from our main specification.

To investigate business stealing, intensive margin, and spillover effects, we resort to a traditional 2SLS model, since each of these questions involve multiple endogenous variables and estimating MTEs with multiple endogenous variables in a single equation is beyond

Figure A3: Alternative Payment Indicators



	(1)	(2)
ATE	0.271 (0.0686)	0.165 (0.0718)
LATE	0.337 (0.0718)	0.255 (0.0813)
F-stat	21.50	20.99

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 two specifications (1–2). The shaded gray indicates 95% confidence intervals.

the scope of this paper. 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 Col 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.

We next explore the role of business stealing between the two branded manufacturers in column 3 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 drugs (and the outside good). But the nature of this substitution is constrained 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 drug (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 drug'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 payment relationships 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. The effect of rival firm meals on own-drug prescribing is estimated to be negative (-0.040), which would indicate more business stealing than suggested by the nested logit functional form. However, neither of these changes are statistically significant, or meaningfully different relative to the baseline specification.

Lastly, we explore the possibility of spillovers between physicians within the same hospital (or practice). Recent work by Agha and Zeltzer (2019) found such spillovers to be important for the diffusion of new drugs. Column (4) of Table A10 reports the 2SLS results where we include the share of a physician's fellow cardiologists in the same hospital or medical practice that have meal-based relationships with the focal firm. In this specification, the point estimate on the main effect of the physician's own meal receipt is similar to all other specifications, although estimated less precisely. The point estimate of the coefficient on "Share colleagues..." equals 0.26. However, our coefficients are imprecisely estimated, so we hesitate to draw strong conclusions about the presence of spillovers in this setting. This lack of precision is driven by the fact that our identification approach assumes all physicians in the same hospital or practice are impacted similarly by spillovers from AMCs in terms of

their meal propensities; i.e., we have little independent variation to identify the effects of "own meal" and "colleagues' meal spillovers."

Table A10: Extensive/Intensive Margins, Business Stealing, and Spillover Tests

	(1)	(2)	(3)	(4)
Meal, own firm	0.396 (0.116)	0.438 (0.122)	0.456 (0.086)	0.317 (0.220)
Meal \$-value, own firm	0.0005 (0.0016)	-0.0006 (0.0013)		
Meal, other firm			-0.040 (0.138)	
Share colleagues w/ meals, own firm				0.257 (0.264)
\$-value demean	Full-sample	Firm-specific	n/a	n/a

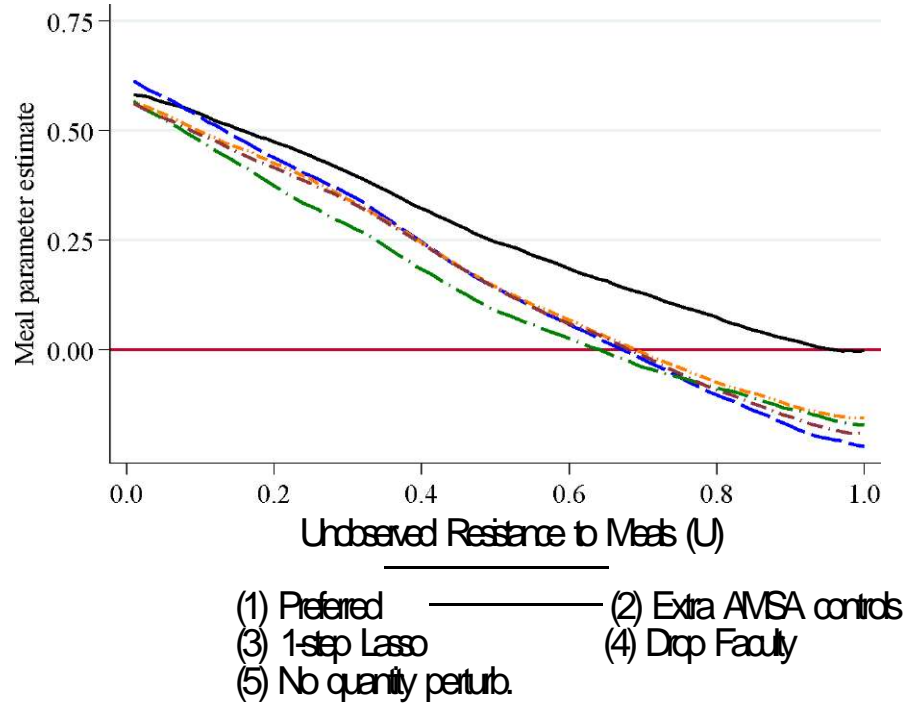
Notes: Replicates the 2SLS-version of the results reported in Figure 4 with additional endogenous variables. "Share colleagues w/ meals" is a variable that indicates the share of a cardiologists' fellow physicians (also cardiologists) within their hospital or practice that also have meal-based interactions with the focal firm.

G.6 Alternative MTE Results and Specifications

Figure A4 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 using the penalty that minimizes the cross-validated MSE in the variable selection routines (the preferred model uses a "two-step" Lasso regression and a penalty that is the largest penalty tested that yields a cross-validated MSE within one standard error of the MSE-minimizing penalty). Specification (4) excludes all AMC faculty from the entire estimation routing (the preferred model includes them). Specification (5) does not perturb the annual claim quantities at the use-case level (the preferred model does).

H Additional Counterfactual Results and Robustness

Figure A4: Alternative MTE Specifications



	(1)	(2)	(3)	(4)	(5)
ATE	0.271 (0.0686)	0.165 (0.0718)	0.175 (0.0635)	0.141 (0.0869)	0.164 (0.0708)
LATE	0.337 (0.0718)	0.255 (0.0813)	0.249 (0.0757)	0.218 (0.0857)	0.241 (0.0783)

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 five specifications (1-5).

Table A11: Payment and Pricing Distortions Table, 2012

	Observed	Ban, fix p	Ban	Ban, $p = mc$	No Ban, $p = mc$
$Q_{statins}$	0.188 (0.001)	0.178 (0.002)	0.178 (0.002)	0.188 (0.002)	0.202 (0.001)
$Q_{Lipitor}$	0.060 (0.001)	0.050 (0.002)	0.050 (0.002)	0.058 (0.002)	0.068 (0.001)
$q_{Crestor}$	0.024 (0.000)	0.016 (0.001)	0.017 (0.001)	0.024 (0.002)	0.036 (0.000)
$oos_{statins}$	17.56 (0.21)	17.56 (0.21)	17.38 (0.21)	2.15 (0.1)	2.15 (0.1)
$POS_{os_{statins}}$	63.88 (0.71)	63.88 (0.71)	63.15 (0.73)	80.85 (1.68)	82.74 (1.58)

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. "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. $N = 124,876$ doctor-drug-brand-year observations with standard errors clustered at the doctor level ($N_d = 15,063$) via delete-120 jackknife and state level via delete-7 jackknife.

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

<i>sde</i>	-350	-100	-50	0
Observed, 2011	89.80	44.75	35.74	26.73
	(0.69)	(0.55)	(0.53)	(0.52)
Ban, 2011	-2.73	-0.53	-0.10	0.32
	(0.37)	(0.04)	(0.03)	(0.08)
Observed, 2012	3.31	1.29	0.89	0.48
	(0.10)	(0.09)	(0.09)	(0.10)
Ban, 2012	-0.11	0.61	0.75	0.89
	(0.52)	(0.12)	(0.06)	(0.06)
Observed, 2011	87.65	42.61	33.59	24.57
	(0.66)	(0.55)	(0.53)	(0.49)
Ban, 2011	-2.54	-0.34	0.08	0.52
	(0.29)	(0.04)	(0.10)	(0.17)
Observed, 2012	3.74	1.72	1.31	0.91
	(0.10)	(0.09)	(0.09)	(0.09)
Ban, 2012	0.54	1.25	1.39	1.49
	(0.46)	(0.07)	(0.06)	(0.11)
Observed, 2011	78.96	33.94	24.94	15.94
	(0.65)	(0.55)	(0.53)	(0.52)
Ban, 2011	-0.68	1.48	1.92	2.36
	(0.06)	(0.32)	(0.39)	(0.46)
Observed, 2012	6.01	3.98	3.58	3.17
	(0.12)	(0.11)	(0.10)	(0.10)
Ban, 2012	4.40	5.03	5.16	5.30
	(0.28)	(0.21)	(0.29)	(0.37)
Observed, 2011	10.82	10.82	10.82	10.82
	(0.09)	(0.09)	(0.09)	(0.09)
Ban, 2011	-2.03	-2.03	-2.03	-2.03
	(0.36)	(0.36)	(0.36)	(0.36)
Observed, 2012	-2.68	-2.68	-2.68	-2.68
	(0.04)	(0.04)	(0.04)	(0.04)
Ban, 2012	-4.43	-4.43	-4.43	-4.43
	(0.32)	(0.32)	(0.32)	(0.32)

Notes: Authors' calculations of equilibrium surplus measures, in dollars per cardiovascular patient. For Observed 2012, Ban 2011, and Ban 2012, surplus measures are shown relative to that Observed in 2011. "Meal Ban" counterfactuals allow both prices and quantities to adjust, per supply and demand model described in text. N = 124,876 doctor-drug-brand-year observations with standard errors clustered at the doctor level ($N_d = 15,063$) via delete-120 jackknife and state level via delete-7 jackknife.

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_{-DE}	-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	18.81	18.81	18.81	18.81	18.85	18.41	18.17	17.94
		(0.23)	(0.23)	(0.23)	(0.23)	(0.21)	(0.31)	(0.40)	(0.30)
	Ban	-0.29	-0.29	-0.29	-0.29	0.94	2.14	2.72	3.78
		(0.04)	(0.04)	(0.04)	(0.04)	(0.23)	(0.52)	(0.67)	(1.04)
POS Statins	Observed	77.04	77.04	77.04	77.04	77.34	75.65	74.48	73.34
		(0.72)	(0.72)	(0.72)	(0.72)	(0.64)	(1.07)	(1.59)	(1.28)
	Ban	-1.27	-1.27	-1.27	-1.27	4.11	9.33	11.92	16.49
		(0.17)	(0.17)	(0.17)	(0.17)	(1.01)	(2.28)	(2.87)	(4.55)

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 (left panel: $\alpha^{de} = 0$) and in Appendix C.5 (right panel: $\alpha^{de} = 1$). 2011 only. "Ban" surplus measures are shown relative to that in the Observed scenario. Results shown for select values of e^{de} . $N = 124, 876$ doctor-drug-brand-year observations with standard errors clustered at the doctor level ($N_d = 15,063$) via delete-120 jackknife and state level via delete-7 jackknife.

Table A14: Robustness of Welfare Estimates to Modeling Assumptions

	e^{de}	-350	-100	-50	0
Total Surplus, Ban 2011	Baseline	-2.73 (0.37)	-0.53 (0.04)	-0.10 (0.03)	0.32 (0.08)
	Rebates	-2.72 (0.37)	-0.53 (0.04)	-0.10 (0.03)	0.32 (0.08)
	MC	-2.72 (0.37)	-0.55 (0.04)	-0.12 (0.03)	0.29 (0.08)
	Pricing	-3.03 (0.42)	-0.77 (0.06)	-0.28 (0.03)	0.21 (0.06)
	ϵ_{DE}	-3.89 (1.75)	-0.89 (0.47)	-0.28 (0.23)	0.32 (0.08)
Consumer Surplus, Ban 2011	Baseline	-2.54 (0.29)	-0.34 (0.04)	0.08 (0.10)	0.52 (0.17)
	Rebates	-2.52 (0.29)	-0.34 (0.04)	0.08 (0.10)	0.51 (0.17)
	MC	-2.50 (0.28)	-0.33 (0.04)	0.08 (0.10)	0.52 (0.17)
	Pricing	-2.91 (0.35)	-0.71 (0.03)	-0.28 (0.04)	0.15 (0.11)
	ϵ_{DE}	-3.68 (1.68)	-0.68 (0.44)	-0.09 (0.23)	0.52 (0.17)
Consumer Surplus, Ban 2011 (-Transfers)	Baseline	-0.68 (0.06)	1.48 (0.32)	1.92 (0.39)	2.36 (0.46)
	Rebates	-0.65 (0.06)	1.50 (0.32)	1.93 (0.40)	2.37 (0.46)
	MC	-0.59 (0.06)	1.54 (0.33)	1.97 (0.39)	2.40 (0.46)
	Pricing	-1.32 (0.11)	0.58 (0.17)	0.88 (0.22)	1.13 (0.26)
	ϵ_{DE}	-1.84 (1.48)	1.08 (0.45)	1.73 (0.35)	2.36 (0.46)
Producer Surplus, Ban 2011	Baseline	-2.03 (0.36)	-2.03 (0.36)	-2.03 (0.36)	-2.03 (0.36)
	Rebates	-2.05 (0.37)	-2.05 (0.37)	-2.05 (0.37)	-2.05 (0.37)
	MC	-2.10 (0.37)	-2.10 (0.37)	-2.10 (0.37)	-2.10 (0.37)
	Pricing	-1.70 (0.31)	-1.36 (0.24)	-1.18 (0.21)	-0.92 (0.20)
	ϵ_{de}	-2.03 (0.36)	-2.03 (0.36)	-2.03 (0.36)	-2.03 (0.36)

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 A12), and alternative specifications: "Rebates" (alternative rebates as described in Appendix E); "Marginal Costs" (extreme alternative assumption that $mc = 0$); "Pricing" ($a^{de} = 1$ in model in Appendix C.5 with insurer sophistication, whereas baseline specification sets $a^{de} = 0$); and " ϵ_{DE} " (an alternative specification with $e^{\wedge} = Y^{de} * Q^{\wedge}$ rather than fixed e^{de} across all physicians). For the $e^{de} = Y^{de} * \bar{r}^{TM}$ specifications, the column value of e^{de} is the average across sample physicians, given their average meal responsiveness \bar{r}^{TM} . $N = 124, 876$ doctor-drug-brand-year observations with standard errors clustered at the doctor level ($N_d = 15,063$) via delete-120 jackknife and state level via delete-7 jackknife.