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

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

November 30, 2023

Abstract

We study how firms target and influence expert intermediaries. In our 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 case study of an important class of cardiovascular drugs, we estimate heterogeneous effects of payments on prescribing, with firms targeting highly responsive physicians. We also develop a model of supply and demand, which 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. In our estimated model, consumers are worse off with payments, unless there is substantial under-prescribing due to behavioral or other frictions. In a final exercise, we calibrate such frictions using clinical data and estimate enough under-prescribing such that, in this case study, payments benefit consumers.

*UC-Berkeley, The Haas School & NBER, mgrennan@berkeley.edu

†Harvard Business School, kmyers@hbs.edu

‡UW-Madison & NBER, ashley.swanson@wisc.edu

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

The data used in this paper were generously provided, in part, by Kyruus, Inc. We gratefully acknowledge financial support from the Wharton Dean's Research Fund, Mack Institute, and Public Policy Initiative. Spencer Boyum, Jack Collison, Gi Heung Kim, April Meehl, 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, Manuel Hermosilla, 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 used in our analyses are derived based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at 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 researcher(s) and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.This paper subsumes an earlier working paper titled "Physician-Industry Interactions: Persuasion and Welfare."

1 Introduction

In many markets, consumers obtain expert advice before making a purchase decision. This is especially true when 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. 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, payments are not allocated randomly (Fugh-Berman and Ahari 2007). Third, the welfare effects of such payments will generally depend on the potential for substitution to competing therapies, equilibrium price adjustments,[2] and the potential presence of 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; Dickstein 2017; Inderst and Ottaviani 2012).

In this paper, we address these challenges using detailed data, a new instrumental variables (IV) strategy and empirical specification that allows us to estimate physicians' *heterogeneous* responses to payments from firms, and a simple but flexible structural model of demand and pricing that accounts for the role of payments as well as the possibility of choice frictions. We find that there is wide variation in treatment effects across physicians, and that firms target physicians with higher expected responses to payments. We use the estimated

[1] See, 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).

[2] See Hastings et al. (2017) for a review and recent example of the large theoretical and empirical literature suggesting that advertising can substantially soften price competition among firms.

model, combined with a calibration exercise based on clinical trial results, to explore the equilibrium price, quantity, and welfare impacts of a ban on payments.

We illustrate our approach using a detailed case study of the market for statins, an anti-cholesterol drug, in 2011-12. Statins are one of the largest-selling drug categories in history. (We discuss the generalizability to other drug categories in the Conclusion.) The two branded statins in our study, Pfizer's Lipitor and AstraZeneca's Crestor, were heavily promoted: over 75 percent of prescriptions in our data were written by physicians who received payments from at least one firm. Lipitor and Crestor made up nearly 40 percent of statin prescribing in 2011, though their prices were around seven times those of generic alternatives. Thus, statins provide an important example of a market with firm payments promoting expensive branded drugs in conjunction with market power and other potential frictions.[3]

We construct a flexible structural model to serve as a complete framework for understanding the economic forces in this setting, and to simulate counterfactual equilibria. Firms negotiate prices with insurers and allocate payments as a function of drug, region, 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 prescribing 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 prices: a new product with Lipitor's same molecule became available at a much lower price, and many insurers removed branded Lipitor from their formularies. This allows us to identify the parameters that determine price sensitivity and substitution across statins, which are important because they determine the extent to which any effects of firm payments result in business stealing, substitution from generics to brands, or substitution from the outside good into statin treatment. In addition, this step recovers a set of fixed effects for each physician-molecule pair, which encompass the ways in which drug quality, firm payments efforts, and any decision errors jointly influence physician prescribing conditional on prices.

In the second step, we regress those physician-molecule fixed effects on payments, specifically meals. To obtain the causal effect of the payments on prescribing, we use an instrumental variable based on variation in physicians' exposure to AMCs' conflict of interest (CoI) policies as measured by the American Medical Student Association's (AMSA) CoI scorecard. CoI policies restrict firms' abilities to provide payments to affiliated physicians.

[3]For example, statins, like many other effective health care treatments, are thought to be underutilized for informational and behavioral reasons (Maningat et al., 2013; Tarn et al., 2021).

Such policies spill over to nearby physicians who are unaffiliated with the AMC due to geographic economies of scale in marketing.[4] Our focal IV captures the extent to which unaffiliated physicians are affected by spillovers from the strictness of CoI policies at nearby AMCs. This instrument has several attractive features. There is a strong first-stage effect of CoI spillovers on meal receipt, and the pattern of spillovers matches our motivating theory: spillovers are weaker among physicians who are directly subject to strict CoI policies, and among physicians located geographically farther from AMCs. Our instrument is well-balanced across a large number of potentially important covariates, and that balance improves with the inclusion of controls. Finally, our results are qualitatively and quantitatively robust as we vary the set of covariates, which include rich controls for cardiologists' patient populations, or restrict the sample to matched control and treatment physicians. Each of these features supports our assumption that the IV is exogenous with respect to other potentially important determinants of payments and prescribing, particularly after including controls.

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 the firm in our data.[5] We use our AMSA instrument to estimate the distribution of marginal treatment effects (MTEs; Heckman and Vytlacil 2007) of payments on prescribing. In order to control for physician- and market-level prescribing differences, we allow for a large number of potentially relevant physician, hospital, and market characteristics related to payments and prescribing. 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 covariates and sample splitting to ensure that our estimates are robust to errors in this selection process (Belloni et al., 2017; Chernozhukov et al., 2018).

[4]This "spillovers" identification strategy is similar to that in Hastings et al. (2017), which relies on variation in sales force exposure driven by the characteristics of other nearby investors. See also Waldfogel (2007) for a broader discussion of this genre of IV strategy. We build upon Larkin et al. (2017), who focus on the direct effects of these policies and estimate significant reductions in prescribing of promoted drugs at the institutions that impose them.

[5]Our IV approach diverges from several recent papers on physician-industry interactions that use physician fixed effects to address the issue of endogenous allocation of payments to physicians (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 (introducing selection on gains to treatment in addition to more typical selection on level types), 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 and the potentially heterogeneous treatment effects on those physicians (Heckman et al. 2006).

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 34 percent for the average physician, which is roughly equivalent to the impact of a $28 price decline or 0.37 standard deviations in the prescribing heterogeneity across physicians. However, the effect is 0.71 standard deviations for a physician in the 90^{th} percentile of the treatment effect distribution, while the effect in the 10^{th} percentile is not statistically different from zero. Firms target physicians who: (1) have more positive expected treatment effects, (2) would otherwise prescribe below-average shares of the firms' drugs, and (3) have larger patient panels. Thus, the incremental profit due to meals is large among physicians who receive meals, but for most physicians not targeted, we cannot reject the null that a meal would lead to zero extra revenue.

We next analyze how prices are determined in equilibrium.[6] 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 drug prices in the presence of payments. Our bargaining parameter estimates are intuitive, suggesting that branded firms receive 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, such as 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). In our model, payments could reinforce or counteract such frictions.[7] We simulate the welfare impact of a payment ban for a wide range of decision errors, and we combine our revealed preference 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 beyond what is revealed by the demand estimates alone. First, the equilibrium effect of meals is to increase use of the focal branded statins by 19 percent on average. This is smaller than the average treatment effect of meals on prescribing because it accounts for the effects of business stealing in the

[6]Prescription drug utilization is typically relatively insensitive to price, but demand expansion driven by firm payments could lead to price effects in a bargaining model. In our context, whether payments lead to price increases or decreases will depend on how insurers weigh consumer surplus vs. their own costs in negotiating with drug firms. Negotiated drug price changes can impact welfare in the short run via an impact on total government spending (given the deadweight loss of taxation) (Decarolis et al., 2020), and in the long run if they ultimately result in different out-of-pocket premiums.

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

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, though at high cost to consumers and payers, as the branded drugs promoted by payments are expensive relative to generics. Third, payments lead to a large amount of substitution from generic to branded statins, in addition to inducing statin prescriptions that would not have been filled otherwise.

Considering a range of potential decision errors, we also 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. If, however, decision errors bias revealed willingness to pay for statins downward by a substantial amount (i.e., 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 willingness-to-pay and conservative estimates of the dollar value of life-year gains due to statins, which are drawn from the clinical literature. Our calibrated decision error value falls well within the region where payments increase consumer welfare, implying substantial under-prescribing of statins. The implied welfare impact of a payment ban is substantial, amounting to slightly more than half that of generic atorvastatin, one of the largest generic introductions in history.

In addition to detailed estimates for an important case study, this paper contributes a useful new instrumental variables strategy and a flexible empirical framework for estimating and mapping heterogeneous 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 complements studies of payments from firms to physicians in a range of important drug classes (Agha and Zeltzer 2019; Carey et al. 2020). Our focus on heterogeneous 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)). Our equilibrium model contributes to an evolving literature on the price effects of factors that shift drug utilization (Dafny et al., 2022; Tuncel, 2023). 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 paper proceeds as follows: Section 2 provides background. Section 3 presents the model. Section 4 discusses the econometrics and results. Section 5 presents the counterfactuals. Section 6 concludes and discusses opportunities for future research.

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 physician, hospital, and regional variation in the data, with a particular focus on our research design.

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 54 percent and 27 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.[8] Also,

[8]In contrast, in interactions with primary care providers (PCPs) in 2011-2, Pfizer might have promoted Celebrex, Enbrel, Lipitor, Lyrica, Norvasc, Prevnar, Pristiq, Viagra, and Zyvox, and AstraZeneca might have promoted Synagis, Toprol, Seroquel, Atacand, Nexium, Prilosec, and Symbicort. Consistent with this,

although cardiologists accounted for only about 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 (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 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). Since our approach relies on accurately mapping physicians to specific geographic regions, we use supplementary data from the National Plan and Provider Enumeration System (NPPES) to ensure a high

after Lipitor went off patent, payments from Pfizer to cardiologists sharply declined, but no such trend break was observed for PCPs. Results for all of our main analyses run on PCPs are available by request.

degree of confidence in physicians' practice primary practice locations.[9]

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 non-statin cardiovascular (ATC code = "C") prescribing to generate a proxy for the physician's market size, the total number of patients seen by the cardiologist who might potentially need a statin in a given year.[10]

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 consistent practice location data and with at least 500 Part D cardiovascular prescriptions on average in 2011 and 2012. This is approximately the 5^{th} percentile of prescriptions per physician-year. For the second step demand estimation, we also restrict the sample to cardiologist-statin molecule pairs that have at least two non-zero observations

[9] On this location dimension, we require that one of two conditions holds: (1) the primary location listed in the Physician Compare data is in the same HRR as the primary location listed in supplementary data from NPPES files that we were able to obtain for 2005-2010; or (2) we cannot locate the physician in the NPPES files, which we take as a signal of high probability that the physician began practicing in 2010 or 2011. As shown in Appendix Table A2, this restriction had little effect on sample physicians' average prescribing patterns.

[10] In order to avoid any mechanical effect of statin prescribing on market size, we proxy for market size using non-statin cardiovascular claims for the cardiologist-year, times 1.4. In this way, we incorporate variation across physicians in their cardiovascular patient panel sizes and tendencies to prescribe a drug for a given patient, which determines market size from the physician's perspective. Then, we scale those claims up to represent the physician's potential market for cardiovascular prescribing, using a scaling factor that ensures that the statin market share will be contained in the interval $[0, 1]$. The scaling factor of 1.4 is approximately $\frac{x}{1-x}$, where $x \sim 0.58$ is the maximum statin share of cardiovascular claims across all sample cardiologist-years in the data. As discussed below, our results are similar under an alternative approach that uses all cardiovascular claims as the market definition. They are also similar under an alternative approach that drops non-statin claims falling in the largest cardiovascular drug classes containing branded drugs which might theoretically also be subject to payments, though Crestor and Lipitor account for much larger sales shares.

Table 1: Prescribing Summary Statistics

	Prescription count, mean		Out-of-pocket price ($), mean		Point-of-sale price ($), mean	
	2011	2012	2011	2012	2011	2012
Market size	4,096	4,577				
	Market share					
All statins	0.163	0.167				
Crestor (rosuvastatin)	0.029	0.029	33.8	33.1	137.1	160.3
Lipitor (atorvastatin)	0.040	0.013	32.0	62.5	139.5	164.0
Generic atorvastatin		0.048		9.7		32.5
Other generics (3)	0.103	0.092	4.5	3.8	13.3	10.3

Notes: Summary statistics are based on a total of 100,763 cardiologist-drug-year observations. 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.

(which is required to recover the mean utility parameter). The final sample used in our analyses contains about 12,250 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.

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 had a market size of over 4,000 Medicare prescriptions in the cardiovascular class per year, and statins had a market share of roughly 17 percent. The effect of entry by generic atorvastatin in December 2011 was 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-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.[11]

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).[12] These documents are the basis of our payments data, which were generously shared by

[11] Branded manufacturers are not passive when their drugs lose exclusivity. For example, there is evidence that Pfizer aggressively promoted a copay coupon program around this time (Aitken et al. 2018), and offered larger rebates to insurers after generic atorvastatin entry (Arcidiacono et al. 2013). Copay coupons cannot be used by Medicare Part D enrollees, so we omit them from our analysis. In our supply side estimation in Section 4.3, we allow for higher rebates by Pfizer in 2012 and test robustness of our results to this choice.

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

Table 2: Payments Summary Statistics

Panel (a): Fraction of cardiologists receiving payments, by type

	All types		Meals		Travel, speak, or consult		Research	
	Raw	Weighted	Raw	Weighted	Raw	Weighted	Raw	Weighted
Crestor	0.624	0.769	0.616	0.760	0.014	0.027	0.001	0.001
Lipitor	0.344	0.443	0.322	0.419	0.015	0.026	0.001	0.001
Either	0.692	0.783	0.677	0.766	0.027	0.042	0.002	0.002

Panel (b): Distribution of payment amounts $ if >0, by type

		mean	*p10*	*p50*	*p90*	*p99*
Crestor	Any	439	15	59	177	10,462
	Meals	82	15	54	160	548
Lipitor	Any	324	11	33	148	6,447
	Meals	52	11	25	121	324
Either (+)	Any	557	16	76	246	13,089
	Meals	99	15	66	206	594

Notes: Statistics calculated on 21,642 cardiologist-drug observations. For "weighted" estimates, we use 2011 prescriptions 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.

Kyruus, Inc.[13]

Table 2 summarizes our data on payments from firms to physicians. As shown in Panel A, meals account for nearly all of the payments we observe in our data.[14] 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-

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

[14]Among physician-firm pairs that involved meals in our sample, 17 percent also involved non-meal payments, but among those pairs that did not involve meals, less than one half of a percent involved non-meal payments.

based relationships are highly persistent over time: for the firm-years in our estimation sample, 73 percent of physicians receiving a meal in year t also receive a meal in year $t+1$. Further, there is not a large amount of variation in the dollar amount of meals when outliers are excluded: the 90^{th} percentile of the distribution of meal dollar values across (nonzero) observations at the physician-firm-year level was less than $160. While these dollar values are small, 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.4, 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 instead use an indicator for receipt of *any* type of payment (e.g., meals, consulting, speaking, travel, or research).

2.4 Regional Variation and Conflict-of-Interest Policies

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. Cardiologists in the 90^{th} percentile HRR per prescribing used roughly twice as many branded strong statins as those in the 10^{th} percentile. Likewise, cardiologists in the 90^{th} percentile HRR per payments were roughly three times as likely to receive a meal compared to those in the 10^{th} percentile.

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

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

(a) Market Share (p.p.)

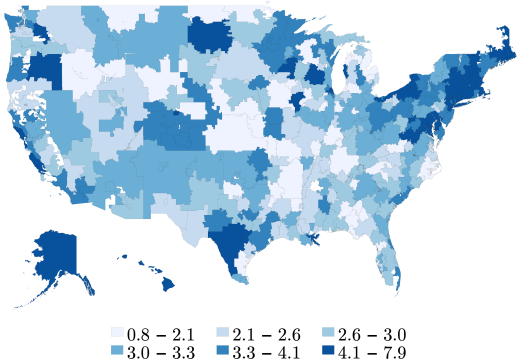

(b) Share Doctors with Meal (p.p.)

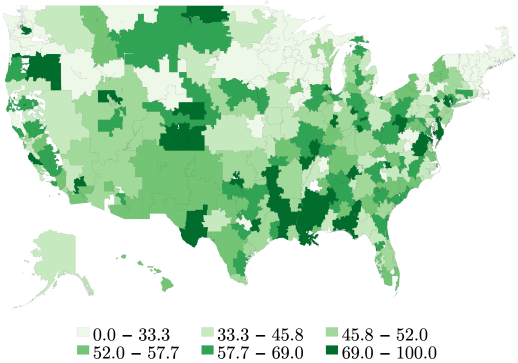

(c) Market Share (p.p.), by Meal Status

		mean	s.d.	*p10*	*p50*	*p90*
Market share of focal drug	Meal	3.32	2.49	0.98	2.70	6.39
	No meal	3.62	2.67	1.10	2.93	6.98

Notes: Panel (a) reports the 2011 HRR-level averages of cardiologist-level market shares for Crestor and Lipitor, averaging over both drugs. Panel (b) reports the HRR-level share of cardiologists receiving meals from AstraZeneca or Pfizer, averaging over both firms. Panel (c) reports the 2011 distribution of cardiologist-level market 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 21,642 doctor-drug level observations from 2011.

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. This pattern may be driven by the prevalence of a number of observed and unobserved factors correlated with prescribing and payments at the region level.

2.4.1 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 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 (Hastings et al. 2017), and 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.[15] Costs include the labor costs of additional sales representatives, the opportunity costs of diverting sales effort from other physicians, and any direct costs of the interaction (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' au-

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

tonomy, pressures on physicians that limit available time for firm interaction, etc. (Khedkar and Sturgis 2015).

We obtain measures of AMCs' conflict of interest policies from the American Medical Student Association's conflict of interest scorecard. We link physicians to AMCs using the Association of American Medical Colleges (AAMC) faculty roster.[16] 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.[17] 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, especially for the cardiologists whose primary practice location is nearer to AMC faculty.

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)). Finally, and most importantly for our purposes, there are spillovers at the regional level, as shown in Panel (c)—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' behaviors. These patterns are consistent with our conversations with current and former pharmaceutical sales executives and pharmaceutical marketing consultants regarding economies of scale in sales force allocation.[18]

[16]Specifically, we link NPIs in Physician Compare to the faculty roster using an algorithm that iterates through the following match criteria, removing at each stage the matches recovered from the previous stage: (1) first name-exact match, last name-exact match, state-exact match, specialty/department-fuzzy match; (2) first name-exact match, last name-exact match, state-exact match; (3) first name-fuzzy match, last name-fuzzy match, state-exact match; (4) first name-exact match, last name-exact match; (5) first name-fuzzy match, last name-fuzzy match. All fuzzy matching is done with the `reclink` package in Stata, using default settings and a match score threshold of 0.995.

[17]In every school year since 2007, medical schools have been asked to submit their policies to the AMSA for rating. Each institution's policy is graded in 13 different categories, including Gifts, Consulting, Speaking, Disclosure, Samples, Purchasing, Sales Reps, On-Campus, Off-Campus, Industry Support, Curriculum, Oversight, and sanctions for Non-Compliance. For each category except Oversight and Non-Compliance, the institution is assigned a numerical value ranging from zero to three. A zero is awarded if the institution did not respond to requests for policies or declined to participate; a one if no policy exists or the policy is unlikely to have an effect; two if the policy represents "good progress" towards a model policy; and a three if the policy is a "model policy." We generate aggregate AMSA scores for each institution by summing across all AMSA components. See Larkin et al. (2017) for more details on the scorecard.

[18]Personal communications: George Chressanthis Jan 5, 2018 and Pratap Khedkar Feb 15, 2018.

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

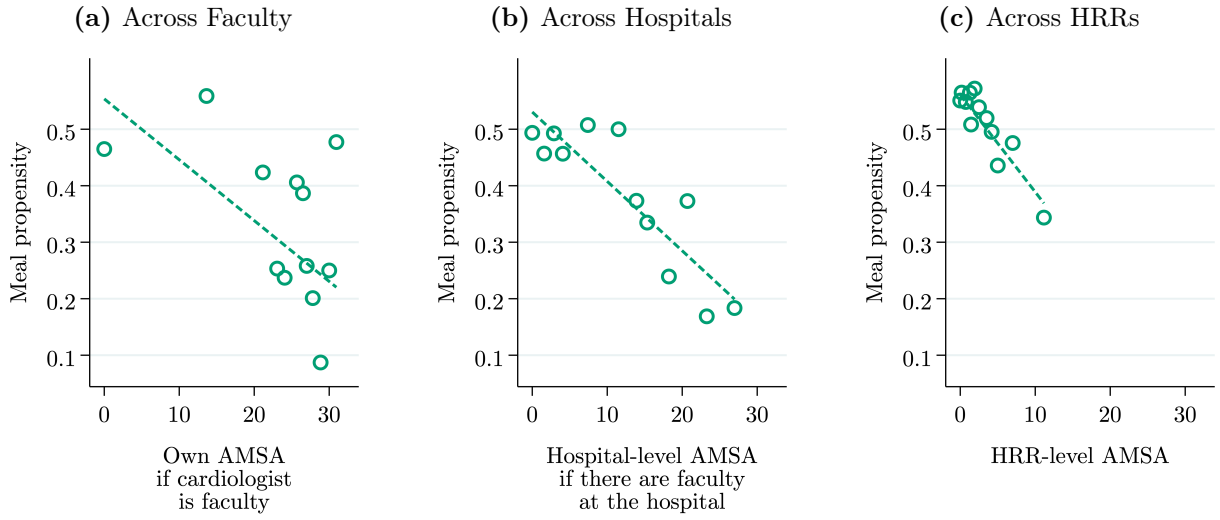

Notes: Panels (a-c) display equally binned scatterplots of the unconditional correlation between meals and three different AMSA score metrics: (a) the cardiologists' own AMC (if they are faculty), (b) the faculty-weighted AMSA score of a cardiologist's hospital (if there is at least one faculty located at the hospital), and (c) the faculty-weighted AMSA score of a cardiologist's HRR excluding the scores of faculty within their own HSA and hospital. The faculty weight is the share of all doctors in the hospital or region that are faculty.

In the remainder of the manuscript, the AMSA score of a cardiologist's HRR, excluding the scores of faculty within their own HSA and hospital and weighted by the share of doctors in the HRR that are faculty (i.e., the variable in Panel (c)), is the instrumental variable we use to identify treatment effects of meals. We include the variables in Panels (a) and (b) as controls.[19]

2.4.2 Physician-Level and Regional Characteristics

The primary concern with using 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).[20] To help ensure that our identifying variation is driven by spillovers from CoI policies rather than these other factors (e.g., preferences,

[19] The geographic distribution of the instrumental variable is shown in Appendix Figure A3.

[20] Another possible concern is that physicians in markets with restrictive policies in our data might have been exposed to more payments in earlier years, prior to changes in AMC policies. We would expect this to attenuate our estimates of the effects of payments on prescribing. It is difficult to assess this directly with our data, but we note that this concern is mitigated if the effects of earlier payments have decayed by the time of our study. Larkin et al. (2017) estimate sharp changes in prescribing of promoted drugs when detailing/meal policies changed at certain AMCs in 2006-12, consistent with such decay.

market structure, etc.), we control for a rich set of observable physician, hospital, and regional characteristics. Here, we outline how we use these controls in our research design 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 CMS' programs for electronic prescribing, electronic health records, or quality reporting. We also include rich data on cardiologists' patient panels from CMS' Medicare Part B public use files. For each sample cardiologist, we include in our set of possible controls: the mean and $10^{th}/25^{th}/75^{th}/90^{th}$ percentiles of the average allowed amount per service for all services provided by the cardiologist, the mean and percentiles of the allowed amount for medical cardiac services, the mean and percentiles of the allowed amount for surgical cardiac services, the mean and percentiles of the allowed amount for emergency services, the service count for each of the above categories of services, the service count associated with new patients, the patient count, and the patient count for medical cardiac care. Each allowed amount is the payment provided for a unit of service, and is the sum of three relative value units (RVUs) times a dollar value. RVUs incorporate regulators' estimates of the intensity and effort associated with different procedures, as well as differences in practice expenses and medical liability insurance associated with different services (Chan and Dickstein, 2019). These variables represent a rich set of controls for the size and severity of the cardiologist's patient panel, as well as for the cardiologists' tendency to provide relatively intensive services.

We supplement this set of cardiologist-specific characteristics with: (1) hospital-level data on admission and bed counts and teaching hospital status from the American Hospital Association, numbers of affiliated physicians across all specialties and cardiologists specifically and affiliated cardiologists' market sizes per the CMS files, and the share of affiliated cardiologists that are faculty and AMSA CoI scores from our own crosswalks with the AMSA data; (2) HSA- and HRR-level aggregates of the hospital variables, HSA- and HRR-level Medicare Advantage eligibility and penetration from CMS data, and HSA- and HRR-level measures of uninsurance rates, Medicaid enrollment rates, and cardiac hospitalization rates from the Behavioral Risk Factor Surveillance System; and (3) ZIP code-level measures of local TV advertising for each of the two branded drugs from the Nielsen AdIntel database. Appendix G.1 reports the summary statistics for all controls, along with the results from univariate regressions of our utilization and meal payment variables on each covariate.

Appendix Table A8 presents results on covariate balance for a number of potentially important variables at the individual cardiologist, hospital, and regional levels. Briefly, there are statistically significant but economically small differences between observations

with low and high AMSA spillover instrument values, and these shrink once one controls for the covariates included in our preferred specification below. This is reassuring evidence that our chosen instrument is largely exogenous with respect to other potentially important determinants of payments and prescribing, particularly after including controls. As we later discuss, our estimated treatment effects are also qualitatively and quantitatively robust as we vary the set of included controls or restrict the sample to matched control and treatment physicians.

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 $j \in \mathcal{J} = \{0, 1, ..., J\}$, for use case i (a doctor/patient/visit combination) in each market defined by doctor d in year t, be given by: $u_{idjt} = \delta_{djt} + \varepsilon_{idjt}$.[21] The choice $j = 0$ represents the choice of treatment other than a statin, with mean utility normalized to $\delta_{d0t} = 0$. We measure the market size of potential statin patients 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 $\varepsilon_{idjt} = \epsilon_{idt} + (1 - \lambda)\epsilon_{idjt}$ is the random coefficients representation of the nested logit model (Cardell 1997), where ϵ_{idt} is a random 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 $\lambda \in [0, 1]$ approaches

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

1, there is less substitution outside the nest of statins.[22] This is an important parameter because it will influence the extent to which meal payments and price competition tend to induce business stealing between statins vs. new prescriptions for patients who otherwise would have received no statin. While business stealing can improve allocative efficiency in an oligopoly setting (Inderst and Ottaviani, 2012), it can also distort demand in favor of expensive branded drugs and away from more cost-effective generics, or it can simply be an inefficient waste of resources as in an advertising "prisoner's dilemma."

We specify mean utility across use cases as:

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

Here, $1_{\{m_{dj}>0\}}$ is an indicator for whether cardiologist d received a meal from j's manufacturer and θ^m_{dj} is 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, $\theta^p p^{oop}_{djt}$ is the average out-of-pocket price paid by patients, multiplied by its utility weight. This degree of sensitivity to price plays an important role in determining the extent to which expensive branded drugs will be close substitutes for cheaper generics, and more broadly the extent of market distortion away from the efficient allocation due to oligopoly power of the branded drugs. $X_{djt}\theta^x_j$ 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, ξ_{djt} is a cardiologist-drug-year-level unobservable term, which we allow to have two components:

$$\xi_{djt} = \tilde{\xi}_{djt} + \varepsilon^{de}_{djt} \ . \tag{2}$$

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

[22] In Appendix G.2, 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."

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

Given a set of drugs $\mathcal{J}_{dt}$ 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:

$$q_{djt} = Q_{dt} Pr[u_{idjt} > u_{ikdt}, \forall k \in \mathcal{J}_{dt}] = Q_{dt} \left(\frac{e^{\frac{\delta_{djt}}{1-\lambda}}}{\sum_{k \in \mathcal{J}_{dt}} e^{\frac{\delta_{kdt}}{1-\lambda}}} \right) \left(\frac{\left(\sum_{k \in \mathcal{J}_{dt}} e^{\frac{\delta_{kdt}}{1-\lambda}} \right)^{1-\lambda}}{1 + \left(\sum_{k \in \mathcal{J}_{dt}} e^{\frac{\delta_{kdt}}{1-\lambda}} \right)^{1-\lambda}} \right) . \tag{3}$$

Given this model, we represent expected consumer surplus as:

$$CS_{dt}(\mathcal{J}_t) = \underbrace{Q_{dt} \frac{1}{\theta^p} \ln \left(1 + \left(\sum_{j \in \mathcal{J}_{dt}} e^{\frac{\delta_{djt}}{1-\lambda}} \right)^{1-\lambda} \right)}_{CS \text{ implied by decision utility}} - \underbrace{\sum_{j \in \mathcal{J}_{dt}} q_{djt} \left(\frac{\varepsilon^{de}_{djt} + \theta^m_{dj} 1_{\{m_{dj}>0\}}}{\theta^p} \right)}_{\text{adjustment for "decision errors" and meals}} . \tag{4}$$

This is the standard formula derived by McFadden (1978), with a modification that captures the extent to which any decision error or meal payment effect causes prescribing to be more (vs. less) appropriate, conditional on all other variables. The first term reflects the consumer surplus that would be implied by our demand estimates if decision utility were equivalent to actual utility. The second term adjusts consumer surplus for the presence of a decision error that results in under- ($\varepsilon^{de}_{djt} < 0$) or over-prescribing ($\varepsilon^{de}_{djt} > 0$), as well as the countervailing (or reinforcing) effect of meals.[23][24] See Appendix C.2 for further discussion.

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

[24] In 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

3.2 Targeting Meal Payments to Physicians

The parameter θ^m_{dj} in the demand model describes the effect of a meal interaction on a physician d's use of branded statin j. 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 allocation. This first stage selection equation takes the form of a linear probability model:

$$1\{m_{dj} > 0\} = X_{dj}\gamma^x_j + Z_{dj}\gamma^z_j + \mu_{dj} \ . \tag{5}$$

where again $1\{m_{dj} > 0\}$ is an indicator for whether cardiologist d received a meal from j's manufacturer. In this equation, X_{dj} is a rich set of covariates that proxy for firms' assessment of the doctor-specific return on investment, Z_{dj} is an instrumental variable representing how regional academic medical center policies spill over into costs of access to unaffiliated doctors in the region, and μ_{dj} is an error term. Appendix C.3.1 shows the tight relationship between Equation (5) and a stylized structural version of this model for a particular cost function with increasing returns to scale. In particular, these targeted meal allocation choices are akin to an entry or advertising game where cost is a function of all the meal allocation decisions across doctors in each region. In keeping with our focus on estimating the effect across doctors who are persistently targeted over time, we conceptualize this as a simultaneous game of incomplete information where the competing firm's expected strategy is captured in $X_{dj}\gamma^x_j + Z_{dj}\gamma^z_j$, which potentially includes information about a doctor that is common across firms, plus a term μ_{dj} which captures information that both firms and the econometrician don't observe about the competitor (as well as information about firm j unobserved only to the econometrician). Our focus on the treatment effects of meals and counterfactual meal bans only requires this semi-parametric equation to model selection into meals. We leave the complexities of solving for the equilibrium of the full meal allocation game (simultaneous or with dynamics) as an interesting and challenging topic for future research.

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

arbitrarily correlated with welfare improvements, depending on the correlation between payments and any underlying decision errors.

equilibrium price changes under a meal payment ban. 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).[25]

Let the supplier's profit be: $\pi(p^{pos}_{jrt}) = \sum_{d \in r} q_{djt}(p^{pos}_{jrt}(1 - \tau_{jt}) - mc_{jt})$, where τ_{jt} is the manufacturer rebate as a fraction of the point-of-sale price and mc_{jt} captures the cost of manufacturing and distributing the marginal unit of drug j in dollars.[26] p^{pos}_{jrt} 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 $p^{oop}_{djt} = cs_{djt}p^{pos}_{jrt}$, 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.[27] 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 Part D plans.[28]

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 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

[25] As discussed by Starc and Swanson (2020), both pharmacies and pharmaceutical manufacturers have market power, but relative market power of different suppliers varies by drug. 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.

[26] When we calculate the gain-from-trade associated with having a given drug on-formulary, we hold fixed the direct and sales force costs associated with payments. This is consistent with the "all else equal" approach of the Nash-in-Nash solution concept we employ below, and with payments being made at the drug firm-physician level rather than the drug firm-physician-insurer level. As a result, those costs fall out of the Nash-in-Nash first-order condition for drug prices.

[27] As discussed in greater detail in Appendix A.2, Crestor and Lipitor typically received "preferred brand" cost-sharing in our baseline setting with meals, so it seems realistic to hold cost-sharing fixed in the counterfactual world without payments, where insurers might be predisposed to treat Crestor/Lipitor more favorably.

[28] Standalone 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.

equations that can be represented by:

$$p^{pos}_{jrt}(1-\tau_{jt}) = mc_j + b_{jrt}\left[\left(1+\sum_{d\in r}\frac{\partial q_{djt}}{\partial p^{oop}_{djt}}\frac{p^{oop}_{djt}-mc_j}{\sum_{d\in r}q_{djt}}\right)\frac{\sum_{d\in r}\widetilde{CS}_{dt}(\mathcal{J}_{dt})-\widetilde{CS}_{dt}(\mathcal{J}_{dt}\setminus j)}{\sum_{d\in r}q_{djt}} + p^{pos}_{jrt}(1-\tau_{jt}) - mc_j\right] \tag{6}$$

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: $\widetilde{CS}_{dt}(\mathcal{J}_{dt}) - \widetilde{CS}_{dt}(\mathcal{J}_{dt}\setminus j)$. Notice that quantities and thus elasticities are driven by physician/enrollee decision-making based on out-of-pocket price p^{oop} given insurance coverage, but the insurer and supplier negotiate over point-of-sale price p^{pos}. The $\widetilde{CS}$ 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 $\alpha^{cs} \in [0,\infty)$ to capture the relative weight insurers place on enrollee surplus and plan costs:[29]

$$\widetilde{CS}_{dt}(\mathcal{J}_{dt}) \;:=\; \underbrace{\alpha^{cs}\,Q_{dt}\frac{1}{\theta^p}\ln\left(1+\left(\sum_{j\in\mathcal{J}_{dt}}e^{\frac{\delta_{djt}}{1-\lambda}}\right)^{1-\lambda}\right)}_{CS \text{ implied by decision utility}} - \underbrace{\sum_j q_{djt}(p^{pos}_{jrt}(1-\tau_{jt})-p^{oop}_{djt})}_{\text{insurer drug costs}}\,.$$

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.[30]

[29] In 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.

[30] The case where insurers perfectly adjust consumer surplus for decision errors and meals introduces some interesting economics, suggesting research to better understand insurer strategies on this dimension could be important. If insurers do adjust in this way, rebates and/or formularies might be more likely to change with a meal ban. Given that Lipitor and Crestor already obtain favorable formulary placement, this seems unlikely to be a first-order change here, but other drugs might be different, reinforcing the importance of further model development and data collection (Dafny et al., 2022; Olssen and Demirer, 2019) on that process.

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 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.[31] 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 AMSA instrument 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: $\delta_{djt} = \ln(s_{djt}/s_{d0t}) - \lambda \ln(s_{dj|_g t})$. Combining this with Equation (1) yields the linear specification:

$$\ln(s_{djt}/s_{d0t}) = \lambda \ln(s_{dj|_g t}) - \theta^p p^{oop}_{djt} + \theta^m_{dj} 1_{\{m_{dj}>0\}} + X_{djt}\theta^x_j + \xi_{djt} \ . \tag{7}$$

where s_{djt} is j's overall market share, s_{d0t} is the market share of the outside good (non-statin treatments), and $s_{dj|_g t}$ is j's market share within nest g, the set of statin treatments. 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 step of estimation, we implement a differences-in-differences style estimator, leveraging the price and choice set variation resulting from the introduction of generic ator-

[31] Appendix D provides full step-by-step details on our algorithm.

vastatin at the end of 2011 to identify the coefficients on price and within-nest share. We estimate:

$$\ln(s_{djt}/s_{d0t}) = \lambda \ln(s_{dj|_g t}) - \theta^p p^{oop}_{djt} + \psi_{dj} + \theta_t + \theta_{\text{Lip12}} + \xi_{djt} \tag{8}$$

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

We account for the endogeneity of $\ln(s_{dj|_g t})$ and p^{oop}_{djt} by constructing instrumental variables that leverage the changes in prices and choice set sizes induced by generic atorvastatin entry. These changes between 2011 and 2012 varied across drug-doctor combinations due to the fact that different insurers prevalent in different regions of the country changed their formularies in different ways, and also the fact that these changes sometimes had different effects on prices and availability of the other strong statin, Crestor, vs. the generic statins from the earlier generation (described in detail in Appendix A.2).[32] Specifically, 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 longer. Also, some insurers moved Crestor to less favorable tiers and/or decreased the number of earlier generation generics on their preferred drug lists. This, combined with different insurer market shares across regions of the country, all resulted in meaningful variation in the relative prices and choice sets consumers faced. To utilize this variation across insurers, while still being careful not to include variation that might be correlated with unobserved changes in regional preferences, we create the following instruments: For each plan-drug-year-region, we take the average out-of-pocket price for that insurer-drug-year across *other* regions, and we then average across plans, weighting by enrollment, to create an instrument for physician d's region: $p^{oop,IV}_{djt}$. We also create an analogous instrument based on an average dummy for formulary inclusion: $\bar{1}_{\{j \in form^{IV}_{djt}\}}$. We denote these two instruments: $Z^p = [p^{oop,IV}_{djt}, \bar{1}_{\{j \in form^{IV}_{djt}\}}]$. Finally, we follow the insight from the literature on instrumental variable choice in nested logit demand systems (e.g., Berry and Waldfogel 1999; Gandhi and Houde 2016) that more variety

[32] An 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 also address this source of bias.

mechanically affects within group shares, by adding polynomials in the cardinality of the sets of statins and strong statins prescribed $Z^J = [\ln(|\mathcal{J}_{dt}|), |\mathcal{J}_{dt}|, |\mathcal{J}_{dt}|^2, \ln(|\mathcal{J}^{ss}_{dt}|), |\mathcal{J}^{ss}_{dt}|, |\mathcal{J}^{ss}_{dt}|^2]$ as an additional set of instruments. These can leverage changes in choice sets closer to the physician level that are not captured by insurer formularies. Our exclusion restriction assumption for identification of (λ, θ^p) is that, within physician-drug, our instruments (Z^p, Z^J) affect the demand dependent variable only through their effects on $(\ln(s_{dj|_g t}), p^{oop}_{djt})$.

4.1.2 Estimating the Effects of Meals on Prescribing

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

$$\hat{\psi}_{dj} = \theta^m_{dj} 1_{\{m_{dj}>0\}} + \theta_j + \bar{X}_{dj}\bar{\theta}^x_j + \bar{\xi}_{dj} \ . \tag{9}$$

The idea of a secondary regression to uncover the determinants of fixed effects goes back at least to Mundlak (1978). The fixed effects are measured with noise, so we employ a version of the standard shrinkage approaches from the empirical Bayes literature.[33] In our preferred specification, we construct $1_{\{m_{dj}>0\}}$ as a dummy for physician d receiving any payment from Pfizer over 2010-2012 (in the case of j=Lipitor/atorvastatin), or as a dummy for physician d receiving any payment from AstraZeneca over 2011–2012 (in the case of j=Crestor).[34] 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. As noted in Section 2.4.1, 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. In the context of our model, this implies the potential for physician selection into treatment on both levels and gains, such that the unobservable in the selection equation μ_{dj} may be correlated with both $\bar{\xi}_{dj}$ and the heterogeneous component of θ^m_{dj}. Accordingly, we develop an econometric specification that

[33] 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 ψ_{dj} distribution.

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

estimates the full distribution of treatment effects θ_{dj}^m across physicians. We estimate the distribution of marginal treatment effects (MTEs) using the `mtefe` package in Stata 16 (Andresen 2018). We then incorporate the estimated θ_{dj}^m's into our counterfactual simulations and report summary treatment effects of interest as a function of the MTEs.

We instrument for $1_{\{m_{dj}>0\}}$ in the above equation using Z_{dj}^{AMSA}, the AMSA score of a cardiologist's HRR, excluding the scores of faculty within their own HSA and hospital, and weighted by the share of doctors in the HRR that are faculty. The exclusion restriction is that, within drug, the instrument Z_{dj}^{AMSA} affects residual physician-drug-level preferences $\hat{\psi}_{dj}$ only through its effect on meals.

The details of the MTE estimation procedure are in Appendix D.2. Identification of θ_{dj}^m requires the standard IV assumptions of relevance, exclusion, and monotonicity, as well as an important assumption, common to most MTE models, of additive separability between observed and unobserved heterogeneity in treatment effects. Intuitively, the level of the treatment effect of meals on prescribing may depend on observables X, but the slope of the MTE curve may not. As noted by Brinch et al. (2017), this assumption is restrictive, but is weaker than and implied by the assumption of additive separability between observables X and the treatment variable, which is standard in applied work using IV.

Our AMSA-based instrumental variables strategy is designed to recover estimates of MTEs as our continuous instrument varies (Heckman and Vytlacil, 2007). Intuitively, if doctors were randomly assigned to regions with either "strict" or "lax" CoI policies, the difference in outcomes across doctors in different regions could be used to measure treatment effects for compliers, the doctors induced to receive meals due to regional CoI policies. Those estimates would be local average treatment effects. The MTEs we estimate using the continuous variation in our instrument are the limits of the LATE parameter as the difference in instrument values becomes infinitesimally small; i.e., as we compare doctors in regions with closer CoI policies. To implement this strategy, the `mtefe` package follows Heckman and Vytlacil (2007) in specifying the conditional expectation of the error term $\bar{\xi}_{dj}$ as a nonparametric function of the treatment propensity given controls $\bar{X}_{dj}$ and the instrument Z_{dj}. It then estimates the conditional expectations of the outcome variable $\hat{\psi}_{dj}$ in the sample of treated and untreated physicians separately, and estimates the MTEs as the pointwise difference between those conditional expectations (Andresen, 2018).

The cross-sectional nature of our identification strategy makes it especially important that we control as well as possible for observable predictors of prescribing patterns and meals. Hence, we include the large set of potential controls at the regional, hospital, and doctor level discussed in Section 2.4.2. In many contexts, a high-dimensional set of possible controls introduces issues with sparsity and collinearity that have been the topic of a growing

literature at the intersection of econometrics and machine learning. We account for this issue by following Belloni et al. (2017), Chernozhukov et al. (2018), and related literature in using Lasso regressions to select the controls which most strongly predict meals and prescribing. In our main results reported in the text, we allow the Lasso to select from among 129 controls, including market size, the physician, hospital, and market variables reported in Tables A5-A7 as well as state fixed effects. In Appendix G.2, we show that results are similar across a range of decisions regarding the set of potential controls, including when we expand the control set to include log and quadratic transformations to allow for more flexible functional forms regarding the relationship between covariates and prescribing/meals.

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 250 bootstrap iterations. Within each iteration, we first drop a random sample of $2 \times \sqrt{N_{\text{d}}}$ 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 250 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 250 estimates. Appendix D presents the estimation routine in full detail, and also presents results on the most frequently selected 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 ($\approx \sqrt{50}$) states in each of the 250 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.2 provides further details, including the robustness of these results to alternative specifications.) Focusing on our main model in the final column, the nesting

Table 3: Demand Estimates Step 1—Price and Nest Coefficients and $\{\psi\}$

	(1)	(2)	(3)
θ^p	0.00173	–0.00027	–0.00816
	(0.00004)	(0.00003)	(0.00018)
λ	0.943	0.961	0.339
	(0.0005)	(0.0009)	(0.0129)
mean(η^p)	0.462	–0.106	–0.209
s.d.(η^p)	0.520	0.119	0.222
N obs.	100,763	100,763	100,763
FE level	d	dj	dj
IV			Y
F stat.			390.6
mean$(\psi_{dj}/\lvert\theta^p\rvert)_{\text{strong statins}}$			–306.7
mean$(\psi_{dj}/\lvert\theta^p\rvert)_{\text{other generics}}$			–327.1
s.d.$(\psi_{dj}/\lvert\theta^p\rvert)$			80.4

Notes: Reports parameter estimates from Eq. 8. Standard errors for the main parameters (θ^p and λ), in parentheses, are based on the perturbed-bootstrap routine described in the text.

parameter estimate of 0.34 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 $\eta^p = \frac{\partial s}{\partial p}\frac{p}{s}$ of -0.21 is within the range of prior estimates for the Part D setting (e.g., Abaluck et al. 2018; Einav et al. 2018). On average, cardiologists value the strong statins about \$20 more than the generics, which is in line with the observed OOP prices (in 2011, the branded strong statins' OOP was around \$28 more than the generics' OOP). The overall physician-molecule preference variation itself is large, with one standard deviation of the ψ_{dj} distribution equivalent to an OOP price differential of about \$80.

4.2.2 Meal Payments First Stage

To explore the first stage effect of spillovers from CoI policies on meal payments, we regress the meal payment indicator on our AMSA instrument: the AMSA score of a cardiologist's HRR, excluding the scores of faculty within their own HSA and hospital, and weighted by the share of doctors in the HRR that are faculty. Figure 3 Panel (a) demonstrates that a standard deviation (2.3 point) increase in our instrumental variable is associated with a 7 percentage point decrease in the likelihood of meal receipt. This first stage relationship is quite powerful, with a median F-statistic across bootstrap iterations of 761. We provide further results regarding this first stage relationship in Appendix C.3.2. Consistent with our motivating theory of AMC policies and geographic spillovers, we document that the

Figure 3: Explaining Variation in Meal Payments

(a) First Stage Regression

	(1)
Standardized HRR-level AMSA IV	–0.0301
	(0.0011)
F stat.	761.2
$L(\mathbf{X})$	Y
N obs.	21,127

(b) Predicted Meal Probabilities

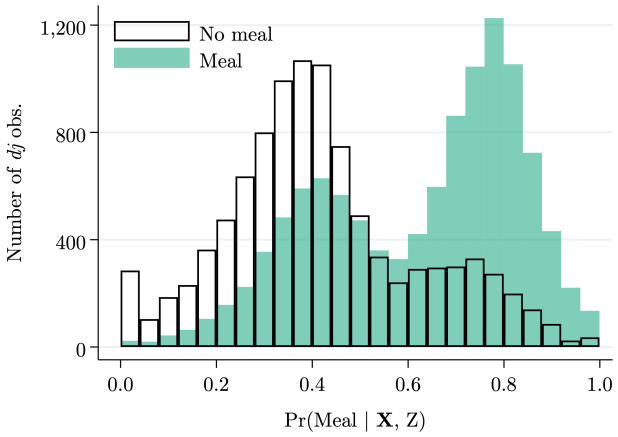

Notes: Panel (a) reports the results from OLS regressions of a meal indicator on the vector of Lasso-selected controls, $L(\mathbf{X})$, and the HRR-level AMSA instrument; the instrument is standardized so that the coefficient indicates the percentage point change in predicted meal probability given a one s.d. change in the IV. Panel (b) displays the distribution of predicted meal probabilities, split by actual treatment status.

first stage is weaker for physicians whose observable characteristics imply they are directly subject to strict CoI policies, which makes them very unlikely to receive a meal payment. Furthermore, we find a weaker first stage among physicians who are geographically farther from AMC faculty.

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. Appendix D.3 provides more details on the controls selected most frequently in the Lasso.

4.2.3 Marginal Treatment Effects of Meals on Prescribing

Figure 4, Panel (a) plots our MTE estimates vs. the unobserved resistance to treatment.[35] The average treatment effect of 0.22 is roughly equivalent to the effect of a $27 price decline, but we reject the hypothesis of a homogeneous treatment effect. At the 10^{th} percentile of unobserved resistance (i.e., physicians that firms are very likely to pay), the effect is approximately 0.5 (equivalent to a $60 price decline), while at the 90^{th} percentile of unobserved

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

Figure 4: Marginal Treatment Effects of Meal Payments

(a) Treatment Effect Distribution

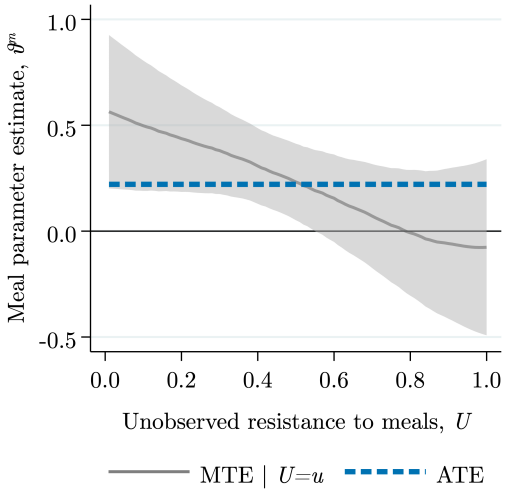

(b) Treatment Effect Point Estimates

		MTE based		
	OLS	ATUT	ATE	ATT
	(1)	(2)	(3)	(4)
θ^m	0.114	0.075	0.221	0.344
	(0.0018)	(0.129)	(0.094)	(0.099)
$L(\mathbf{X})$	Y		Y	
IV			Y	
N obs.	21,127		21,127	

Notes: Panel (a) plots the MTE curve ($E[\theta^m \,|\, U = u]$) with 95 percent C.I. in shaded grey, as well as the ATE estimate for reference. Panel (b) reports the estimates from the OLS (Col. 1) and MTE (Cols. 2–4) specifications.

resistance (i.e., physicians that firms appear to avoid), the point estimate is approximately –0.05 and it is not statistically distinguishable from zero. Appendix G.3 shows that the level and slope of these MTE estimates are similar under alternative approaches to the selection of controls, alternative sampling restrictions, alternative definitions of market size and the payment variable, and alternative Lasso and MTE estimation decisions.

The table in Panel (b) compares several estimates of θ^m: ordinary least squares (column (1)), as well as the average treatment effect on the untreated (column (2)), the average treatment effect overall (column (3)), and the average treatment effect on the treated (column (4)) associated with the marginal treatment effects.[36] There are two key points of interest in this table. First, the estimated average treatment effect of 0.221 is larger than the OLS estimate of 0.114. As discussed in Section 2.4.1 above, we expect that firms would target physicians with larger appropriate patient populations and larger expected treatment effects. The first form of targeting could generate positive bias in the OLS estimate, if our controls were not sufficient to capture the appropriateness of targeted doctors' patient populations. However, targeting on treatment effects would be expected to push the bias in the opposite

[36]Each of these estimates is based on the MTE estimates rather than more conventional estimators like 2SLS; as shown in Heckman et al. (2006), standard instrumental variables methods are not guaranteed to weight marginal treatment effects positively in a setting with selection on gains or "essential heterogeneity." The bias in standard IV can be extreme; e.g., IV can be negative even though all pointwise MTEs are positive.

direction. This logic is outlined in greater detail in Appendix C.3, and Appendix D.2 relates our econometric specification to a generalized Roy model with essential heterogeneity. The fact that the ATE estimate shown in Figure 4 Panel (b) is larger than OLS suggests that the latter effect dominates in this setting. This is consistent with the comparison of the ATT and ATUT, which indicates that physicians receiving meals have larger treatment effects, as predicted in the generalized Roy framework (see Heckman and Vytlacil (2007) for a review). Second, our estimated average treatment effects are larger than those found in other papers that address physician selection into receiving payments with the inclusion of physician fixed effects (Agha and Zeltzer 2019; Carey et al. 2020; Shapiro 2018a). For example, Shapiro (2018a) finds that a detailing visit increases prescribing of antipsychotics by 14 percent in the subsequent year,[37] whereas the coefficients in our nested logit demand model imply that a meal-based relationship increases promoted statin prescribing by about 32 percent for the average physician, but by 54 percent for the average physician actually targeted by firms and only 8 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.[38]

Next, we explore the targeting patterns in greater depth. 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 $E[\theta^m_{dj}]$ for any observation in the data.[39] Figure 5, Panel (a) presents a histogram of expected treatment effects, normalized by the standard deviation of the physician-molecule preference variation, $E[\theta^m_{dj}]/SD(\psi^m_{dj})$, 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 0.60 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 0.08, and for roughly 95 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 (a), which implied a sizable difference between the average treatment effect on the treated and the average treatment effect on the untreated.

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

[38] Chintagunta and Manchanda (2004); Shapiro (2018a); Agha and Zeltzer (2019) each consider the role of detailing "stock." Agha and Zeltzer (2019) also explicitly focus on diffusion of drugs at the beginning of their life cycles.

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

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

(a) Prescribing Change with Payment (s.d.)

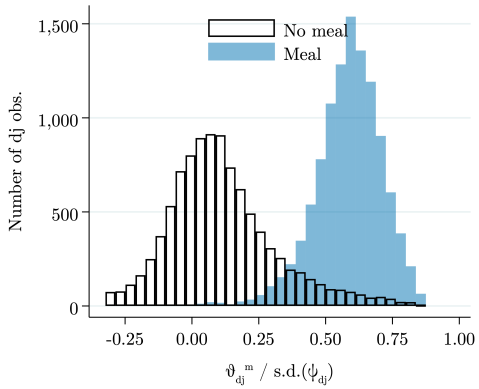

	p10	*p25*	*p50*	*p75*	*p90*
No meal	–0.12	–0.02	0.08	0.21	0.38
Meal	0.43	0.52	0.60	0.67	0.74

(b) Prescribing Without Payments (share)

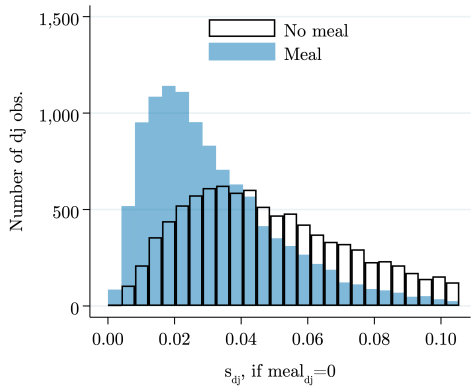

	p10	*p25*	*p50*	*p75*	*p90*
No meal	0.02	0.03	0.05	0.07	0.10
Meal	0.01	0.02	0.03	0.04	0.06

(c) Manufacturer Return to Payment ($)

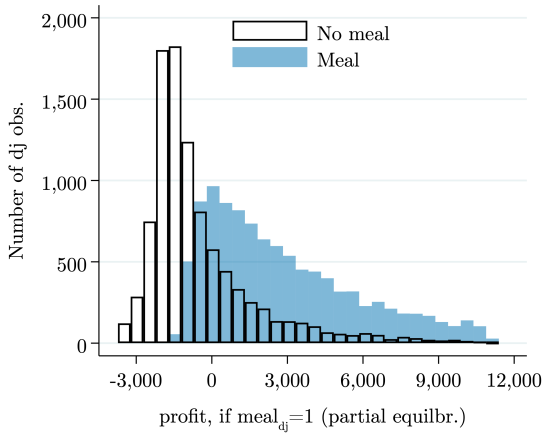

	p10	*p25*	*p50*	*p75*	*p90*
No meal	–2,394	–1,862	–1,207	171	2,806
Meal	–364	518	2,417	5,593	10,078

Notes: Panel (a): expected meal response parameters $E[\theta_{dj}^m]$, scaled by the standard deviation of ψ_{dj}; Panel (b): expected cardiologist-branded drug-specific market shares, setting each meal indicator to zero $E[s_{djt}^{m=0}]$; and Panel (c) the distribution of manufacturers' incremental profits due to meals $(p_{jrt}^{pos}(1-\tau_{jt}-w_{jt})-mc_{rt}) \times E[q_{dj}^{m=1}-q_{dj}^{m=0}]-C_j$, 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 τ, and markups charged by supply intermediaries w (see Section 4.3 below and Appendix E for details). Tables below each plot report select percentiles of these distributions by treatment status. All computations for 2011 only.

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, $E[s_{dj}(m_{dj} = 0)]$. 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. In our profit calculation, we account for two distinct costs: the dollar value of the meal payments and the average sales force cost associated with physician-firm interactions.[40] The distribution for treated physicians is shifted significantly rightward from that of untreated physicians. Meals increased profits to drug firms by roughly $2,417 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 -$1,207. These estimates provide insight regarding why some physicians are targeted by firms and others are not.

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 (α^{cs}), the decision error (ε^{de}), the manufacturer rebates (τ_{jt}), and the

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

marginal costs (mc_{jt}).

To estimate the model for a given vector $(\varepsilon^{de}, \tau_{jt}, mc_{jt})$, we parameterize bargaining ability parameters as a function of drug and regional fixed effects, and specify the econometric unobservable as the residual variation in bargaining parameters needed to fit the model to the data. The literature on empirical models of business-to-business bargaining has identified a potential endogeneity problem caused by the fact that prices can affect the surplus created, and thus for any particular negotiation, the econometric unobservable in the pricing equation may be correlated with the observable surplus measures going to the buyer and supplier. To address this, we follow the literature (Grennan 2013; Brown 2019) in using the consumer surplus measures for each drug-region, calculated at average prices for the same drug in other regions, as instruments. The logic behind this strategy is that these are correlated with the variation in the surplus measures coming from the demand estimates of product qualities and substitution parameters, but uncorrelated with the focal drug-region pricing unobservable.[41] We then estimate the insurer weight on consumer surplus and bargaining ability parameters via GMM.

We still need to address the challenges of estimating manufacturer rebates (τ_{jt}) and the marginal costs (mc_{jt}). 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: $mc = 0.17 * \overline{p^{pos}_{gen}}$. 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 (τ_{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

[41] In practice, this simultaneity problem will be quite small to the extent that the pricing unobservable is a small determinant of price, and/or price is a small determinant of the surplus created. With drug-time specific bargaining parameter fixed effects in the pricing equation, and with the combination of cost sharing and low consumer price sensitivity in demand, one could argue that is the case here.

Table 4: Supply Parameter Estimates

		Atorvastatin	Lipitor	Lovastatin	Pravastatin	Crestor	Simvastatin
$b_{j,2011}$			0.550	0.047	0.043	0.552	0.036
			(0.043)	(0.005)	(0.004)	(0.041)	(0.004)
$b_{j,2012}$		0.119	0.659	0.033	0.033	0.640	0.026
		(0.013)	(0.031)	(0.004)	(0.004)	(0.049)	(0.003)
α^{cs}	2.326						
	(0.117)						

Notes: $N = 105,865$ doctor-drug-brand-year observations with standard errors clustered at the doctor level ($N_d = 11,870$) via delete-220 jackknife and state level via delete-7 jackknife.

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.[42]

Finally, we estimate that the weight insurers place on enrollee surplus in negotiations is larger than the weight they place on net costs: $\alpha^{cs} = 2.33$. 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.[43]

5 Equilibrium Welfare Effects of Meals

The above results demonstrate that meals have heterogeneous effects on prescribing, and that they are targeted to more 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

[42]Our pricing model should still be sufficient to predict counterfactual prices as long as one does not think a meal ban would have substantially changed how this process of Lipitor going generic unfolded in 2012. In the welfare analysis in the next section, we focus on the counterfactual estimates for 2011, using 2012 primarily as another way to calibrate the magnitude of the meal ban effects in 2011 to the magnitude of the generic atorvastatin entry effects.

[43]We do not model such subsidies—e.g., risk corridors and reinsurance—because they are applied at to insurers' overall enrollee population rather than at the drug or drug class level.

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 $p^{oop} = mc$ is efficient at one extreme where $\varepsilon^{de} \geq 0$, and no ban and $p^{oop} = mc$ is efficient if ε^{de} is negative and large enough. Table 5 displays several key results from these counterfactuals for 2011. 2012 results are qualitatively similar and shown in Appendix Table A10.

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

	Observed	Ban, fix p	Ban	Ban, $p = mc$	No Ban, $p = mc$
$Q_{statins}$	0.157	0.149	0.150	0.163	0.174
	(0.001)	(0.002)	(0.002)	(0.002)	(0.001)
$Q_{Lipitor}$	0.034	0.029	0.029	0.040	0.047
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
$Q_{Crestor}$	0.022	0.016	0.016	0.023	0.031
	(0.000)	(0.001)	(0.001)	(0.002)	(0.000)
$OOP_{statins}$	18.73	18.73	18.53	2.15	2.15
	(0.21)	(0.21)	(0.22)	(0.1)	(0.1)
$POS_{statins}$	74.91	74.91	74.10	86.15	87.65
	(0.74)	(0.74)	(0.73)	(1.43)	(1.32)

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. $N = 105,865$ doctor-drug-brand-year observations with standard errors clustered at the doctor level ($N_d = 11,870$) via delete-220 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 market share by 0.7 percentage points—a five percent reduction in total statin usage. For the focal branded statins, the decrease is 0.12 percentage points, which is a 19 percent reduction on average.

These numbers tell us several important things about how meal payments induce substitution toward the focal drugs: (1) Meals induce enough substitution among those patients who would otherwise receive no statin to increase total statin usage by five percent. (2) That this increase in total statin usage is smaller than the increase in branded statin usage indicates that meals also induce substitution from the generic statins to the branded statins (about $(0.12 - 0.07)/0.12 = 42$ percent of the increase in branded statin usage comes from this channel). (3) Finally, the increase in branded statin usage is smaller than the meal treatment effects documented in Section 4.2.3 because this counterfactual accounts for the business stealing between the two branded firms among the physicians who receive meals from both.

Turning to the remaining columns of the Table, the quantity estimates also show that pricing above marginal cost reduces total statin usage by about 1.7 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 undershoot the efficient allocation; in the Crestor case, meals cause utilization to fall almost exactly at the allocation with $p = mc$ and a meal ban.

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. The exception to this is shown in the last two columns, which 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 by 15-17 percent.

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 ε^{de} that dictates the extent of under- or over-prescribing of statins relative to the outside option; we explore alternative specifications of the decision error in Appendix H. Our counterfactuals can speak to this issue as they allow for more or less substitution to the outside good (though they do 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 ε^{de}

In our welfare simulations, we present two different measures of consumer surplus: $CS_{dt}(\mathcal{J}_t)$ accounts for surplus net of out-of-pocket prices. $CS_{dt}(\mathcal{J}_t) - \sum_j q_{djt}(p^{pos}_{jrt}(1-\tau_{jt}) - p^{oop}_{djt})$ (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 and/or taxes.

We compute Producer Surplus as the marginal profit as defined in Section 3.3, minus sales force and meal costs: $PS_{jrt} = \sum_{d \in r} q_{djt}(p^{pos}_{jrt}(1-\tau_{jt}) - mc_{jt}) - C_{jr}^{m_{dj}=1}$, where τ_{jt} is the manufacturer rebate, mc_{jt} captures the cost of manufacturing and distributing the marginal unit of drug j, and $C_{jr}^{m_{dj}=1}$ is an estimate of the average sales force and meal costs associated with a meal-based relationship with a physician.[44] Total Surplus is the sum of Producer Surplus and Consumer Surplus net of transfers. We calculate surplus for the "Observed" and "Ban" scenarios in 2011. We also calculate surplus for the "Observed" scenario in 2012, which provides some context for the magnitude of meal effects in that we can compare them to the welfare impact of generic atorvastatin entry.

[44] As discussed previously, we assume that marginal manufacturing costs are 17 percent of the average POS price of generic statins, and that sales force 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.

The results are summarized graphically in Figure 6 with all measures represented as 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 2011" to "Observed 2011" to quantify the welfare effect of a meal ban in 2011. For context, Appendix Table A11 provides estimates of Observed 2011 levels: per patient total surplus was \$21.52, with \$12.22 accruing to consumers (net of transfers) and \$9.27 accruing to producers (assuming $\varepsilon^{de} = 0$).[45] The measures are represented in dollars per cardiovascular patient in order to capture the full welfare impact of the counterfactual policy changes (in particular, this measure properly takes into account changes on the extensive margin of patients who move in/out of receiving a statin under different policies). Some readers may find it useful to multiply these by six (which is roughly the inverse of the average market share of statins) to compare these numbers to average prices of statins, for example.

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 ε^{de}. Intuitively, the more negative ε^{de} is, the greater the implied benefit of taking statins, and in turn of the statin market expansion in 2012. At $\varepsilon^{de} = 0$, Consumer Surplus increased by \$1.97 per patient due to generic atorvastatin entry, but Consumer Surplus net of transfers increased even more (\$3.85) due to the latter measure incorporating the full benefit of reduced POS prices.

In contrast, Consumer Surplus (with or without transfers) under a meal ban is rotated counter-clockwise, relative to Observed Consumer Surplus. More negative ε^{de} implies that statins are more valuable to patients, and hence that a meal ban has more potential to be harmful. The point at which the line for "Observed 2011" crosses the line for "Meal Ban 2011" 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 $\varepsilon^{de}/\theta^p$ <-\$46. However, meals only improves consumer surplus *net of transfers* for the more extreme threshold value of $\varepsilon^{de}/\theta^p$ <-\$225, reflecting that underutilization must be more extreme for the market expansion effect of meals on overall statin use (which is smaller than the effect of meals on use of promoted statins) to be valuable enough to justify increased insurer costs for 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

[45] Appendix Table A11 also provides the numbers underlying the changes from these levels shown in Figure 6 (along with Meal Ban results specific to 2012), and standard errors for all estimates.

Figure 6: Welfare and Counterfactual Estimates

(a) Consumer Surplus

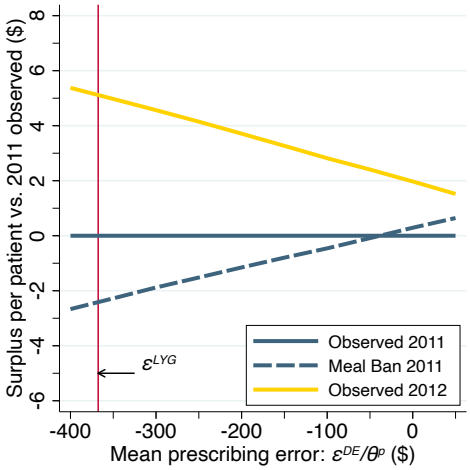

(b) Consumer Surplus – Transfers

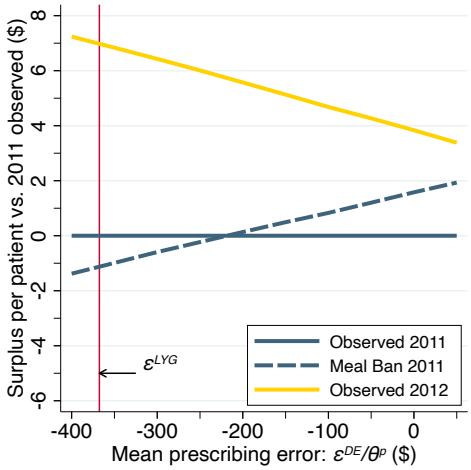

(c) Producer Surplus

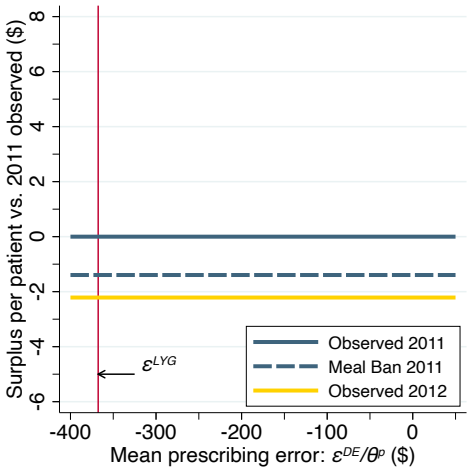

(d) Total Surplus

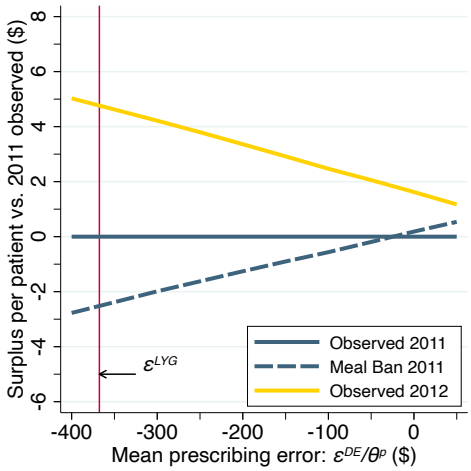

Notes: Authors' calculations of equilibrium surplus measures, in dollars per cardiovascular patient, relative to that Observed in 2011. "Meal Ban" counterfactuals allow both prices and quantities to adjust, per supply and demand model described in text. Results shown for $\varepsilon^{de} \in [-\$400, \$50]$. Detailed results for select values of ε^{de} available in Appendix Table A11.

we estimate here is fairly large, with a meal payment ban resulting in an approximately 15 percent decrease in producer surplus.

Taken together with Consumer Surplus net of transfers, Panel (d) shows that, in the case of statins, meals increase Total Surplus as long as $\varepsilon^{de}/\theta^{p}$ <-\$18. If there is no underlying decision error, a meal ban increases Total Surplus by \$0.20 per patient. Alternatively, if decision errors are equivalent to the average effect of meals on revealed willingness to pay ($\varepsilon^{de}/\theta^{p} = ATE(\theta^{m})/\theta^{p}$=-\$28), 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 A12 shows how our welfare simulations vary with our modeling assumptions, comparing the above ("Baseline") results to simulation results with alternative assumptions. For each alternative specification, we show the effect of a meal ban on 2011 surplus (in dollars per cardiovascular patient) for a range of possible values of ε^{de}. In one, we re-compute the counterfactuals with mean price elasticities of -0.36 and -0.06, one standard deviation higher and lower than our baseline estimate of -0.21, based on the variation across drugs observed in Einav et al. (2018) (which also covers the range between long- and short-run elasticities for statins measured in Feng (2022)). We consider this to be an interesting robustness check because our identification strategy relies on relatively short-run changes to statin prices and choice sets, and our sample of cardiologists could be more likely than the average physician to be writing patients' first statin prescriptions. Thus, one might think our estimates measure something between the short- and long-run price elasticities for statins in the population we study. From this perspective, this exercise provides a robustness check regarding potential effects in the long- vs. short-run or for drugs with different elasticities, showing that even this range is still informative for evaluating meal ban policies, as the conclusions are ultimately similar.

Besides elasticities, we also consider robustness to assumptions regarding: rebates, marginal costs, the extent to which insurers internalize cardiologist decision errors ("Pricing"), and decision errors being correlated with meal responsiveness $\varepsilon_{d}^{de} = \gamma^{de} * \bar{\theta}_{d}^{m}$ instead of fixed across all physicians. The effects of a meal ban are qualitatively and quantitatively similar across all modeling assumptions. There are a few interesting patterns where modeling decisions can interact with meal ban policies in interesting ways, though: 1) If insurers understand and incorporate decision errors into price negotiations, then a meal ban can lead to an increase in prices of the branded drugs, thereby dampening the effect of the meal ban on producer surplus. 2) If doctor under-prescribing is correlated with receipt of a meal, then the wel-

fare loss from a meal ban accrues more quickly with the overall level of under-prescribing. Overall, these findings indicate that, in addition to the degree of under/over-prescribing, optimal policy toward meals will depend on other factors like the sophistication of insurers, how meals are allocated, and consumer price sensitivity.

5.2.2 Calibrating the Decision Error Magnitude using Clinical Data

As a final exercise, we incorporate external evidence on the likely magnitude of decision errors into the above welfare framework. The extent of over- and underutilization (absent meals) surely varies across drugs. Many studies of statins point to potential underutilization due to patients not clearly understanding the benefits and/or wanting to try alternatives such as diet and exercise (Tarn et al., 2021), physicians providing insufficient rationale/information to patients, and physicians not conducting the necessary lipid testing to diagnose high cholesterol in the first place (Maningat et al., 2013). These all fit in the "behavioral hazard" model of Baicker et al. (2015), which includes a decision error that accommodates potential mechanisms such as lack of information, overoptimism, present bias, salience, and forgetfulness that might underpin the underuse of statins.[46] This perspective is also consistent with our result that meals tend to bring otherwise low prescribers closer to the prescribing behavior of those who do not receive meals. However, determining whether ε^{de} is sufficiently negative 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 (around $33, in the case of branded statins Crestor and Lipitor) to a "flow" willingness-to-pay value of $516.[47] This measure suggests that the $\varepsilon^{de}/\theta^{p}$ consistent with statin demand in our data

[46] Statins are far from the only example of underutilization of proven therapies in health care. Baicker et al. (2015) cite the underuse of glucose-lowering drugs for diabetes, beta blockers for heart disease, medication to control HIV, antibiotics for tuberculosis, prenatal care, and immunosuppressants after organ transplants. And these are but a few of the examples from the medical literature on under-adherence relative to evidence-based medicine (van Dulmen et al., 2007). Importantly, several studies have shown that small changes to financial incentives such as copayments (Choudhry et al., 2011; Brot-Goldberg et al., 2017) and overall drug budgets (Chandra et al., 2023) have large impacts on subsequent patient health outcomes, consistent with health benefits being dramatically underweighted in decision utility.

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

would be -$368 (indicated by a vertical red line in each Panel of Figure 6). This is well below even the 2011 threshold value of -$225 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 61 percent of (randomly selected) Medicare patients in our sample would experience the clinical benefits of statins from the medical literature, even if the remaining 39 percent experienced zero benefit. And so on.

We find this flexibility an appealing feature of the "decision error" approach to modeling frictions between decision utility and welfare relevant utility in health care. One can use a relatively transparent set of assumptions to map clinical data to revealed preference demand estimates for a given sample. For illustrative purposes, we have done so at the aggregate 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 (as we do) 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 a little more than half that associated with the introduction of generic atorvastatin, one of the largest generic introductions of all time.

Our counterfactual simulations also show how the overall effect of meals in the statin

market comes through a combination of inducing new statin prescriptions among patients who otherwise would not have received a statin, inducing switching from inexpensive generics to more expensive branded statins, and business stealing between the two branded drugs often promoted via meals to the same doctors. This makes it clear that the structure of the market—in terms of the number of branded drugs, viability of generic competitors, and preferences governing price sensitivity and substitution to outside treatments—all play important roles in determining the welfare effects of meal payments.

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 is a useful starting point for those explorations.

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.

In particular, one interesting and useful area to explore is the dynamics of payments and interactions—how they accumulate and decay over time and in turn how this determines short- and long-run effects on prices and prescribing. Understanding such details could have important implications for manufacturer promotion strategies as well as public policy, and how those might vary across different drugs with different market structures and information environments.

A broader study that measures responses of prescribing to payments across many drugs would also be 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.

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.

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.

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.

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

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

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

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

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

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

Bhattacharya, V., Illanes, G., and Padi, M. (2020). Fiduciary duty and the market for financial advice. Working paper.

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

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

Brot-Goldberg, Z. C., Chandra, A., Handel, B. R., and Kolstad, J. T. (2017). What does a deductible do? the impact of cost-sharing on health care prices, quantities, and spending dynamics. *The Quarterly Journal of Economics*.

Brown, Z. Y. (2019). An empirical model of price transparency and markups in health care. *working paper*.

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, D. C. and Dickstein, M. J. (2019). Industry input in policy making: Evidence from medicare. *Quarterly Journal of Economics*, 134(3):1299–1342.

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., Flack, E., and Obermeyer, Z. (2023). The health costs of cost-sharing. *NBER Working Paper 28439*.

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.

Choudhry, N. K., Avorn, J., Glynn, R. J., Antman, E. M., Schneeweiss, S., Toscano, M., Reisman, L., Fernandes, J., Spettell, C., Lee, J. L., Levin, R., Brennan, T., and Shrank, W. H. (2011). Full coverage for preventive medications after myocardial infarction. *New England Journal of Medicine*.

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.

Dafny, L., Ho, K., and Kong, E. (2022). How do copayment coupons affect branded drug prices and quantities purchased? *accepted in American Economic Journal: Economic Policy*.

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.

Decarolis, F., Polyakova, M., and Ryan, S. P. (2020). Subsidy design in privately provided social insurance: Lessons from medicare part d. *Journal of Political Economy*, 128(5).

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).

Feng, J. (2022). History-dependence in drug demand: Identification and implications for entry incentives. *The Review of Economics and Statistics*.

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

Gandhi, A. and Houde, J.-F. (2016). Measuring substitution patterns in differentiated products industries. Working paper.

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.

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).

Maningat, P., Gordon, B., and Breslow, J. (2013). How do we improve patient compliance and adherence to long-term statin therapy? *Curr Atheroscler Rep.*, 15.

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.

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

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

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

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

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

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

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

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

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

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

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

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

Tarn, D. M., Pletcher, M. J., Tosqui, R., Fernandez, A., hong Tseng, C., Moriconi, R., Bell, D. S., Barrientos, M., Turner, J. A., and Schwartz, J. B. (2021). Primary nonadherence to statin medications: Survey of patient perspectives. *Preventive Medicine Reports*, 22.

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

Tuncel, T. (2023). Should off-label drug prescriptions be prevented? empirical evidence from france. *accepted in Review of Economic Studies*.

van Dulmen, S., Sluijs, E., van Dijk, L., de Ridder, D., Heerdink, R., and Bensing, J. (2007). Patient adherence to medical treatment: a review of reviews. *BMC Health Services Reviews*.

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

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

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

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

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

APPENDICES—FOR ONLINE PUBLICATION ONLY

A Additional Institutional Background

A.1 Medicare Part D

37 million people, or 70 percent of eligible Medicare beneficiaries, enrolled in Part D plans in 2014 (Hoadley et al. 2014). Medicare-eligible individuals can acquire prescription drug coverage through standalone Part D plans or bundled with medical and hospital coverage in the form of "Medicare Advantage" plans. Utilization of drugs in the Part D program 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.[48]

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.[49]

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

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

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

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 formulary tiers, point-of-sale price, and benefit phase. A typical formulary structure would have generic drugs on tier 1, preferred branded drugs on tier 2, and non-preferred branded drugs on tier 3, though finer distinctions between drugs are possible.[50] In the Part D setting in 2011, the generic drugs lovastatin, pravastatin, and simvastatin had median formulary placement on tier 1, and both Crestor and Lipitor had median formulary placement on tier 2. 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.[51]

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.[52] 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.[53]

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

[50] https://www.medicare.gov/drug-coverage-part-d/what-medicare-part-d-drug-plans-cover

[51] In 2011, the maximum out-of-pocket price for LIS beneficiaries with income above 100 percent of the federal poverty level (FPL) was $2.50 for generic drugs and $6.30 for branded drugs, and many LIS beneficiaries qualified for more generous subsidies based on income.

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

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

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 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.	β^{cross}	s.e.$^{\text{cross}}$	mean	s.d.	β^{cross}	s.e.$^{\text{cross}}$	β^{panel}	s.e.$^{\text{panel}}$
Lipitor	OOP	31.74	6.07	0.77	0.25	66.58	14.64	0.82	0.12	1.02	0.04
	POS	139.64	2.02	0.94	0.96	163.12	8.31	1.70	1.09	0.96	0.05
Atorvastatin	OOP					9.78	1.57	0.29	0.14		
	POS					31.43	3.27	0.75	0.22		
Crestor	OOP	32.29	6.37	0.75	0.30	31.90	5.70	0.34	0.28	0.80	0.44
	POS	137.34	2.09	0.62	1.24	160.54	1.79	0.62	0.32	1.00	0.01

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 (β reports the "first stage" regression coefficient with the standard errors in the next column). There is a strong positive association between the pricing policies of the dominant insurers in each state and their pricing policies in other regions. This holds within each year, looking across states cross-sectionally ("cross"), and within states, looking across years, which we can see in the final "panel" column that pools years and controls for state fixed effects. These associations are generally more precise for OOP prices (which we use in our demand analysis) than for POS prices. This suggests that the correlation in "insurer-specific factors" across regions is stronger for benefit design (e.g., formulary structure) than for 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. Column (1) includes all cardiologists in CMS' Physician Compare database. Column (2) includes the subset of those physicians matched to a "reliable" practice location. A practice location is defined as reliable if either: (a) We found the physician's NPI in CMS' National Plan and Provider Enumeration System (NPPES), and the NPPES and Physician Compare addresses were located in the same HRR; or (b) we did not find their NPI in the NPPES. Column (3) includes the subset of column (2) physicians with at least 500 cardiovascular prescriptions in the Part D data on average in 2011 and 2012 and non-zero usage of a given molecule in multiple years (for estimation purposes). Column (4) includes the subset of those physicians with positive prescribing of Crestor or Lipitor in both 2011 and 2012. Column (5) is a superset of column (3) for use in our counterfactual analyses, including observations from cardiologists with non-zero usage of a molecule only in one year.

In terms of the two main regressions used to identify the demand parameters: the step 1 demand regression, which estimates the price and nest parameters, is based on data at

the doctor-molecule-brand-year level (*djbt*; Panel a) for all drugs and uses the sample corresponding to column (3); the step 2 demand regression, which estimates the distribution of meal parameters, is based on data at the doctor-molecule level (*dj*; Panel b) for Crestor and Lipitor only, and uses the sample corresponding to column (4); the counterfactuals that incorporate the supply-side are all based on data at the doctor-molecule-brand-year level (*djbt*; Panel a), and use the sample corresponding to column (5).

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 (NPI).

We construct two main categories of payments: "research" and "general" (all non-research payments). This scheme closely follows that of Open Payments and excludes all royalty payments. Within general payments we identify three sub-categories: "meals," "travel or lodging," and "consulting, speaking or education."

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 $\overline{Q^{eff}Q^{b}}$—marginal patients prescribed the drug in the presence of payments to physicians receive negative health benefits.

In a setting with perfect competition, this conceptual framework suggests that 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

Table A2: Sample Descriptions

		(1)	(2)	(3)	(4)	(5)
Panel (a): Cardiologist-drug-brand-time (*djbt*, all drugs) – Claims data						
Cardiologists	unique	19,296	14,839	12,237	11,718	12,248
Market size	total	1.3e+08	1.1e+08	1.1e+08	1.0e+08	1.1e+08
	mean	3,708	4,027	4,337	4,456	4,334
	median	2,912	3,277	3,568	3,695	3,565
Market share, all statins	mean	0.181	0.178	0.165	0.168	0.168
	median	0.166	0.166	0.161	0.163	0.163
Market share, Crestor	mean	0.033	0.032	0.029	0.029	0.028
	median	0.024	0.023	0.023	0.023	0.022
Market share, Lipitor-2011	mean	0.047	0.046	0.040	0.040	0.040
	median	0.035	0.035	0.034	0.034	0.034
Market share, atorvastatin-2012	mean	0.056	0.053	0.048	0.048	0.047
	median	0.044	0.043	0.042	0.042	0.042
Market share, other generic sum	mean	0.109	0.105	0.098	0.098	0.099
	median	0.097	0.096	0.093	0.094	0.095
N *djbt* obs.	unique	141,603	115,760	100,763	99,249	105,865
Panel (b): Cardiologist-drug (*dj*, Crestor or Lipitor) – Payment data						
Cardiologists	unique	16,467	13,438	11,734	11,718	11,870
AstraZeneca – Crestor	Any type	0.709	0.721	0.729	0.729	0.724
	Total $ amount	455	467	438	438	440
	Any meal	0.700	0.711	0.719	0.719	0.714
	Meal $ amount	83.3	84.1	82.4	82.4	81.9
Pfizer – Lipitor	Any type	0.346	0.353	0.354	0.354	0.351
	Total $ amount	299	308	270	270	269
	Any meal	0.323	0.330	0.332	0.333	0.330
	Meal $ amount	53.2	53.5	51.7	51.7	51.6
N *dj* obs.	unique	29,157	24,100	21,146	21,127	21,642
Panel (c): Cardiologist (*d*, Crestor and/or Lipitor) – Payment data						
Cardiologists (N *d* obs.)	unique	16,467	13,438	11,734	11,718	11,870
All Types	either firm	0.657	0.671	0.691	0.692	0.687
	both firms	0.318	0.323	0.322	0.322	0.321
	$ sum	376	403	385	385	381
Meals	either firm	0.641	0.656	0.676	0.677	0.672
	both firms	0.300	0.305	0.304	0.304	0.303
	$ sum	64.2	66.9	67.1	67.1	66.4

Notes: Reports select summary statistics for prescribing- and payment-related outcomes at three levels of observations (Panels (a–c)) and across five samples as described in the text.

Figure A1: Welfare Analysis with Other Frictions

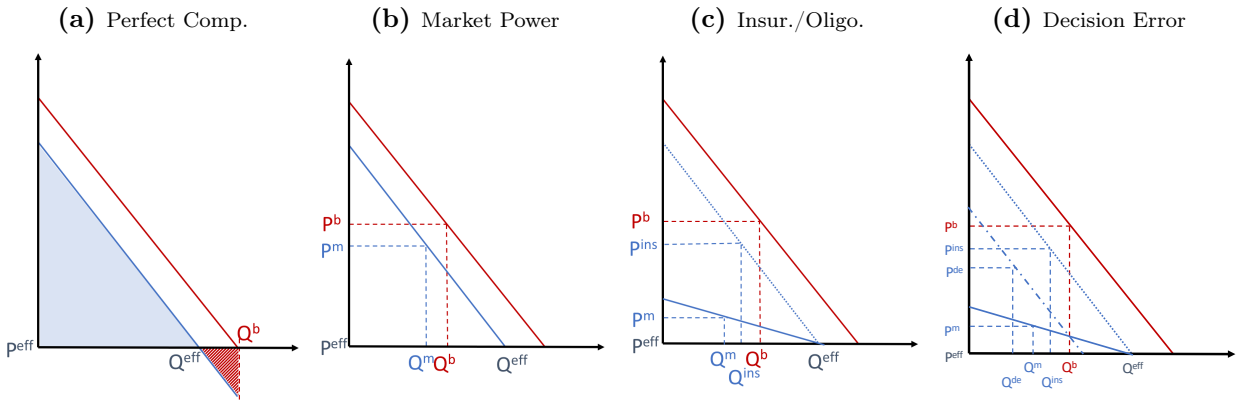

is distorted away from the social optimum due to high prices. In prescription drug markets, branded drugs have patent protection, and they often compete with differentiated branded and generic substitutes whose manufacturers make their own strategic pricing and promotion decisions. Payments are typically 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 $Q^m < Q^b < Q^{eff}$ (pictured) or cause prescribers to overshoot the optimum $Q^m < Q^{eff} < Q^b$. In the former case, the overall welfare impact of payments is positive, though consumer surplus declines; in the latter case, both total and consumer surplus decline.

Finally, we must also account for reasons that the "effective" demand curve for a given drug may not represent the appropriate one for welfare analysis. A leading example is insurance, pictured in Panel (c) of Figure 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: $Q^m < Q^{ins} < Q^b$.[54] The welfare implications are again ambiguous, and the consumer surplus effects of firm payments will depend on pass-through of producer

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

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 under-prescribing 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 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:

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

where $\tilde{\xi}_{djt}$ is a typical demand unobservable that impacts both choices and true realized utility, but ε^{de}_{djt} is a "decision error" that impacts decision utility but not realized utility.

Given this model, we represent expected consumer surplus as:

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

The second term adjusts consumer surplus for the presence of a decision error that results in under- ($\varepsilon^{de}_{djt} < 0$) or over-prescribing ($\varepsilon^{de}_{djt} > 0$) as in Figure 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, $\varepsilon^{de}_{djt} = 0$ would be a case with no decision error at all, where meals simply bias utilization of promoted drugs upward. By contrast, $\varepsilon^{de}_{djt} = -\theta^m_{dj}$ would be a case where meals perfectly correct prescribing errors among those who receive them.

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 $\varepsilon^{de}_{djt} := \varepsilon^{de}$, and we simulate counterfactuals for a range of decision errors, from substantial under-prescribing to over-prescribing.[55] This specification, while simple, has the virtue of being easy to interpret, and accommodates the finding in the prior literature that statins *as a drug class* are under-prescribed (Baicker et al. 2015).

In an alternative specification, we simulate welfare under the assumption that decision errors are a scaled function of estimated physician-specific meal responses $\varepsilon^{de}_{d} := \gamma^{de}\bar{\theta}^m_d$, varying scalar γ^{de} to again allow for a range of potential under- or over-prescribing in the absence of meals. In this specification, $\gamma^{de} = -1$ represents a special case where meal payments perfectly correct for a given physician's average tendency to under/over prescribe, for those drugs for which meals are received. This specification also covers cases where meals distort prescribing away from otherwise optimal behavior $\gamma^{de} = 0$, under-correct $\gamma^{de} < -1$ or overcorrect $\gamma^{de} \in (-1, 0)$ under-prescribing, 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

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

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 sales force 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:

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

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

The institutional details in this setting suggest that the cost function $C_{jr}^{m_{dj}=1}$ will have increasing returns to scale in the sense that the average cost of providing a meal will be decreasing in the total meals provided in a region. We would also expect the cost function to depend on other regional characteristics such as the density of candidate physicians in geographic space. Further, the incremental cost of providing a meal to doctor d is likely to depend on characteristics of that doctor or her employer that affect her willingness to accept a meal.

The expected quantity increase from the meal $E[q_{dj}^{m_{dj}=1} - q_{dj}^{m_{dj}=0}|\mathcal{I}_{dj}]$ will be a function of the expectation of total size of the doctor's patient flow Q_{dt} and the choice probability function as given in Eq. (3). In particular, it will be a function of the expectation of the parameter θ_{dj}^{m} which determines the effect of the meal interaction on the mean utility weight the doctor assigns to drug j.

In keeping with our focus on the estimating the effect across doctors who are persistently targeted over time, we conceptualize this as a simultaneous game of incomplete information where the competing firm's expected strategy is captured by integrating out its effect on the choice probabilities that govern the quantity response at the doctor level, as well as the overall expected price equilibrium. Our focus on meal bans only requires a model selection equation. We leave the complexities of solving for the equilibrium of the full meal allocation game (simultaneous or with dynamics) as an interesting and challenging topic for future

research.

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 of the meal selection equation and variables included in our instrumental variables analysis.

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

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

To deconstruct this expression, we use $\partial q / \partial 1_{\{m>0\}}$ as an approximation to $\left(q_{dj}^{m=1} - q_{dj}^{m=0}\right)$.[56]

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

Substituting these values gives

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

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

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

Here, as discussed above, the rival firms' meal strategy will enter at the doctor level through its expected impact on $\frac{\partial s_{dj}}{\partial 1_{\{m_{dj}>0\}}}$ and at the regional level through expected equilibrium prices $p_{jr_d}^{mfr}$.

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

C.3.2 First Stage Consistency with Motivating Theory

Table A3 sheds additional light on the first stage relationship between our focal instrument, the HRR-level AMSA CoI instrumental variable, and meal receipt. The first column replicates the main result in Figure 3. Columns (2) and (3) interact the AMSA IV with indicators for subsamples of physicians whose observable characteristics imply they are directly subject to strict CoI policies, which makes them very unlikely to receive a meal payment regardless of the attractiveness of the local physician population from the firm's perspective. The results indicate that the first stage relationship is weakened for faculty physicians at high-AMSA (strict CoI) AMCs (column 2) and physicians practicing at hospitals associated with strict CoI policies (column 3). In the notation of our motivating theory above, meals are shut down among physicians with high access cost A_{dj}.

Next, in column (4) we investigate how the first stage relationship varies with physician proximity to AMCs. In our simple model, meals will depend on N_{jr_d}, the number of other doctors accessed in the region near physician d. To explore this intuition, we interact the AMSA IV with the distance from the focal physician to faculty cardiologists. We expect the AMSA IV to reduce meals for a cardiologist more strongly when a cardiologist is geographically close to AMC faculty; the positive coefficient on the faculty distance interaction term in column (4) confirms this hypothesis. This is not simply driven by local physician density: in fact, the negative non-faculty distance interaction term shows that the AMSA IV first stage is weaker when the focal cardiologist is closer to other non-faculty cardiologists. This is consistent with our general returns-to-scale argument, in that firms will be more inclined to engage with cardiologists in denser areas to capitalize on fixed costs of such engagements.

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:

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

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:

Table A3: Alternative First Stage Regressions

	(1)	(2)	(3)	(4)
HRR-level AMSA IV	–0.0301	–0.0312	–0.0339	–0.0395
	(0.0011)	(0.0011)	(0.0011)	(0.0013)
... ×**1**{High own AMSA}		0.0212		
		(0.0052)		
... ×**1**{High hosp. AMSA}			0.0383	
			(0.0023)	
... × dist. from faculty				0.0012
				(0.0003)
... × dist. from non-faculty				–0.0125
				(0.0010)
N obs.	21,127	21,127	21,127	21,127

Notes: Reports results from regressions of the meal indicator on the AMSA instrument, including alternative interaction terms: the cardiologist has above-median own AMSA scores (Col. 2), above-median hospital-level faculty-weighted AMSA scores (Col. 3), or their mean standardized distance from other cardiologists that are, or are not, faculty (Col. 4).

$\widetilde{CS}_t(\mathcal{J}_t) - \widetilde{CS}_t(\mathcal{J}_t \setminus j)$. τ_{jt} reflects the rebate rate to insurers off the posted price p^{pos}_{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 under-prescribing of some drugs and take that into account when negotiating prices. In such a scenario, we would want to replace $\widetilde{CS}_{dt}$ in equation (6) with the following:

$$\widetilde{CS}_{dt}(\mathcal{J}_{dt}) := \alpha^{cs}\Bigg[\underbrace{Q_{dt}\frac{1}{\theta^p}\ln\left(1+\left(\sum_{j\in\mathcal{J}_{dt}} e^{\frac{\delta_{djt}}{1-\lambda}}\right)^{1-\lambda}\right)}_{CS \text{ implied by decision utility}} - \alpha^{de}\underbrace{\sum_{j\in\mathcal{J}_{dt}} q_{djt}\left(\frac{\varepsilon^{de}_{dj}+\theta^m_{dj}1_{\{m_{dj}>0\}}}{\theta^p}\right)}_{\text{adjustment for "decision errors" and meals}}\Bigg] - \underbrace{\sum_j q_{djt}(p^{pos}_{jrt}(1-\tau_{jt})-p^{oop}_{djt})}_{\text{insurer drug costs}}.$$

Here, we include a parameter $\alpha^{de} \in [0,1]$ 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 ($\alpha^{de} = 0$), the fully "sophisticated" case where insurers perfectly adjust

consumer surplus for decision errors and meals ($\alpha^{de} = 1$), 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 A12 and A13. As shown in Appendix Table A12, the value of α^{de} has little impact on the welfare implications of a meal ban. 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 POS prices increasing by \$12 if $\varepsilon^{de} = 0$, but only by \$3 if $\varepsilon^{de} = -350$. In the former case, meals just decrease consumer surplus, and the insurer internalizes that fact and pays substantially lower prices in the presence of meals. In the latter case, meals offset the underutilization of a high-value drug and the insurer internalizes that as well, such that a meal ban has a much smaller effect on negotiated prices.

D Parameter Estimation Routine

The following outline details the steps necessary to recover the demand parameters (θ^p, λ, ψ_{dj}, θ^m) and is followed by more in-depth discussions of the Lasso approach we use (Appendix D.1), how the MTEs are estimated (Appendix D.2), the important variables selected by the Lasso algorithm (Appendix D.3), and the role of the perturbation and shrinkage procedures (Appendix D.4). 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 250 times, dropping the observations for a randomly selected $2 \times \sqrt{N_{\text{d}}}$ doctors (approx. 220 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 $q_{djt} = 50$, this would translate to 50 rows of $c = 1$ use-cases for that djt)
3. Sample with replacement
4. Calculate perturbed quantities $\tilde{q}_{djt} = \sum c_{djt}$

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

5. Estimate Eq. (8) to recover price ($\theta^{p,k}$) and nest (λ^k) parameters, and doctor-molecule fixed effects (ψ^k_{dj})
6. Parameter estimates: for θ^p and λ, point estimates are from estimation on the full sample, standard errors are given by the standard deviation across the 250 samples
7. Shrink each ψ^k_{dj} estimate towards the j-specific mean using the standard deviation of ψ^k_{dj} across the 250 samples as the standard error in the empirical Bayes shrinkage formula

For each k, estimate meal parameters

8. Keep ψ^k_{dj} estimates for Crestor and Lipitor observations

9. Follow the split-sample Lasso approach described below in Section D.1 to select the relevant controls (X)
10. Estimate MTEs for meal receipt, $\theta^{m,k}$ 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 250 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), and is implemented using the `ivlasso` command in Stata (from the `pdslasso` package developed by Achim Ahrens, Christian Hansen, Mark Schaffer, and Thomas Wiemann (see: https://statalasso.github.io). 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, use `ivlasso` (with all options set to their defaults, partialling out the instrument to ensure it is "selected" in every sub-sample) to select the relevant controls (X)
3. Within each sub-sample s, estimate $\theta^{m,k,s}$ (and other MTE parameters) using the variables selected in the opposite sub-sample s'
4. Solve for the k-specific estimate $\theta^{m,k} = (\theta^{m,k,A} + \theta^{m,k,B})/2$
5. Parameter estimate: for θ^m (and all corresponding MTE-based estimates) point estimates and standard errors are given by the median, $\overline{\theta^m} = median(\theta^{m,k})$, and the median deviation, s.e.$(\overline{\theta^m}) = \sqrt{median(\theta^{m,k} - \overline{\theta^m})}$, respectively.

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, and with W and V denoting unobservables in the outcome and treatment equations, respectively:

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

Then we make the two necessary assumptions of conditional independence ($W^d, V \perp Z \mid X$: the error terms in the outcome and treatment equations are orthogonal to the instrument

conditional on the controls) and separability ($\mathbb{E}[W^d \mid V, X] = \mathbb{E}[W^d \mid V]$). Per this model and assumptions, MTEs are then defined as:

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

where U, the unobserved resistance to treatment, is given by the quantiles of V.[57]

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) in our baseline specification, we use a nonparametric local linear function to estimate the control functions in the model (which are related to $\mathbb{E}[W^1 - W^0 \mid U_i = u]$). We show in Figure A5 in Appendix G.4 that our results are similar if we instead use a more flexible second degree (quadratic) local polynomial.

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

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

where P_i is the doctor's propensity score.

D.3 Important Variables

Table A4 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 250 bootstrap samples, each with two split-samples, the total possible number of selections is 500. The Table reports all controls that were selected in more than 80 percent of the subsamples.

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

Table A4: Frequently Lasso-selected Variables

	share of subsamples selected in
HRR: Cardiac hospitalization rate	1.00
HSA: Faculty share	1.00
HRR: Uninsured rate	1.00
Zipcode-drug: Ad volume, duration	1.00
HSA: Cardiac hospitalization rate	1.00
Cardiologist: Part B, service count, new patients	1.00
HSA: Medicare Adv., enrollment	1.00
HRR: Faculty share	1.00
Cardiologist: Market size	1.00
HSA: Cardiologists, sum	1.00
HSA: Uninsured rate	1.00
Zipcode-drug: Ad volume, spending	1.00
Hospital: Num. cardiologists	1.00
HRR: Teaching hospitals, sum	1.00
Hospital: Is teaching hospital	1.00
HSA: Teaching hospitals, sum	1.00
Cardiologist: Part B, patient count, all services	1.00
Cardiologist: Part B, allow amount, emerg., *p75*	0.99
Cardiologist: Part B, allow amount, all services, *p75*	0.99
Hospital: AMSA, faculty-wgt.	0.99
Cardiologist: Own AMSA	0.98
Cardiologist: Miles to nearest AMC	0.98
HRR: Medicaid, enrollment	0.97
Cardiologist: Part B, service count, cardio. surg.	0.96
Cardiologist: Med. school grad. year	0.94
HRR: Cardiologists' market size, mean	0.93
Cardiologist: Part B, service count, emerg.	0.92
Hospital: Cardiologists' market size, mean	0.88
Cardiologist: Part B, allow amount, cardio. surg., *p25*	0.84
Cardiologist: Part B, allow amount, cardio. med., *p90*	0.80

Notes: Reports the share of subsamples the covariate is selected in either the outcome or treatment Lasso regression; showing only variables selected in at least 80 percent of the subsamples.

D.4 Perturbation and Shrinkage

We are concerned that the doctor-molecule mean utility parameters (the ψ_{dj} fixed effects) might be influenced by noise (since we only observe two years of utilization), especially for 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-220 jackknife bootstrap, blocked at the cardiologist level to allow for arbitrary correlations within cardiologist, where we remove 220 physicians (which is twice 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 99 percent of the observations are shrunk by less than 1 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 p_{jrt}^{pos} that suppliers receive, and that insurers pay after cost-sharing is applied, are net of rebates τ. This is an approximation, as we are collapsing a set of bilateral negotiations between upstream and downstream firms into a single negotiation over a unit price, and the "producer surplus" is split between manufacturers, wholesalers, distributors, and pharmacies. Second, Panel (c) of Figure 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 τ *and* "other suppliers' markup" w, so that manufacturer revenue becomes $R(p^{pos}_{jrt}) = \sum_{d \in r} q_{djt} p^{pos}_{jrt}(1 - w - \tau)$. In the supply side estimation, welfare simulations, and simulations of manufacturer revenue, we rely on researchers' estimates of τ and w, and we test the sensitivity of our results to our decisions on how to use these estimates.

To obtain the components τ and w, we rely on multiple sources.[58] CMS has reported average manufacturer rebate percentages overall ($\tau = 0.175$) and for cardiovascular drugs specifically ($\tau = 0.263$) going back to 2014.[59] Arcidiacono et al. (2013) assume $\tau = 0.151$ and estimate that (in the antiulcer drug market) rebates increase to 48.3 percent after branded drugs' patents expire. Similarly, Aitken et al. (2018) suggest that Lipitor rebates increased in 2012. One can also infer w from Yu et al. (2018), as they pulled together aggregate data on profits to PBMs, wholesalers, pharmacies, providers, and insurers. We ignore the profits of PBMs and insurers, as those are "buyers" in our calculation. We also ignore provider profits, as those refer to physician-administered drugs such as chemotherapy and are not relevant for statins. That leaves wholesalers and pharmacies, which are estimated to obtain profits of $0.037 * p^{pos}$ and $0.152 * p^{pos}$, respectively. Thus, the work in Yu et al. (2018) suggests that $w = 0.190$.

For our simulations of manufacturer revenue, the above papers suggest $\tau + w = 0.32$ if statin markups and rebates look like those of the average drug in the US. If statins instead follow other cardiovascular drugs in having relatively high rebates, then $\tau + w = 0.263 + 0.190 = 0.453$ would be more appropriate. For our supply side estimation and counterfactuals, the above papers suggest $\tau = 0.32 - 0.037 - 0.152 = 0.131$ as a lower bound based on patterns observed for a wide range of pharmacy-dispensed drugs (Kakani et al. 2020; Yu et al. 2018) and $\tau = 0.263$ based on cardiovascular drugs only. We use ($\tau = 0.263, \tau + w = 0.453$) in the main text. As a robustness check, we consider ($\tau = 0.131, \tau + w = 0.32$) in Table A2 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

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

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

Figure A2: Heterogeneity in Profit Effects, Alternative Rebate Assumption

(a) Preferred Rebate Assumption ($)

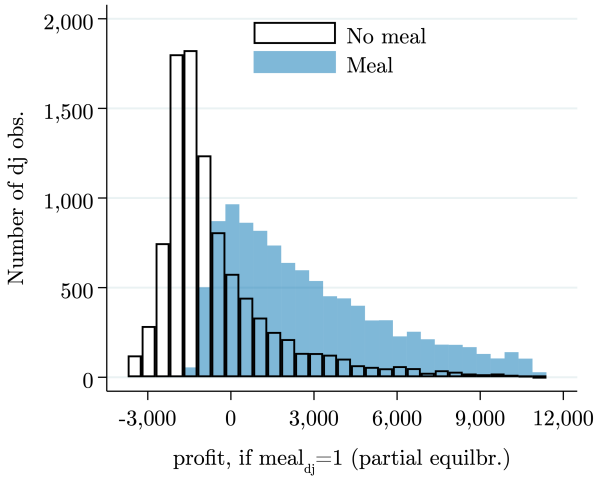

	p10	*p25*	*p50*	*p75*	*p90*
No meal	–2,394	–1,862	–1,207	171	2,806
Meal	–364	518	2,417	5,593	10,078

(b) Alternative Rebate Assumption ($)

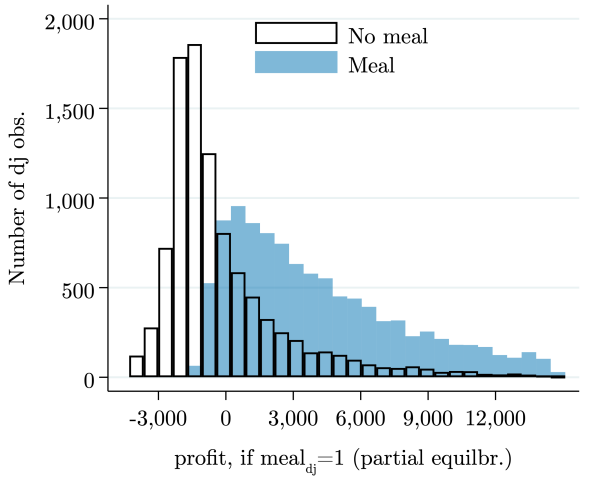

	p10	*p25*	*p50*	*p75*	*p90*
No meal	–2,581	–1,894	–1,060	703	4,094
Meal	–11	1,125	3,571	7,647	13,432

Notes: Plots the posterior (dj-) cardiologist-drug-specific estimates of the distribution of manufacturers' incremental profits due to meals $(p^{pos}_{jrt}(1-\tau_{jt}-w_{jt})-mc_{rt})\times E[q^{m=1}_{dj}-q^{m=0}_{dj}]-C_j$, 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 τ, and markups charged by supply intermediaries w. Panel (a) is our preferred approach, with Panel (b) alternatively assuming $p^{mfr}=0.68*p^{pos}$.

the main text by using $(\tau=0.483, \tau+w=0.673)$ (based on Aitken et al. (2018); Arcidiacono et al. (2013)); we stick with the alternative assumptions $(\tau=0.131, \tau+w=0.32)$ in 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 $(\tau=0.24, \tau+w=0.41)$.

Comparing Table A2 to the results from our preferred specification taken from the main text (Figure 5, Panel c), this alternative (larger) pass-through assumption yields larger profits. 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.

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"; 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 in-

crease 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

Figure A3 presents a map of the geographic distribution of the residual variation in our HRR-level AMSA instrumental variable. AMSA CoI scores are on average higher in the midwest and eastern regions of the country than in the west, but low CoI regions are interspersed fairly evenly throughout the midwest and east.

Figure A3: Residual Variation in AMSA CoI IV across HRRs

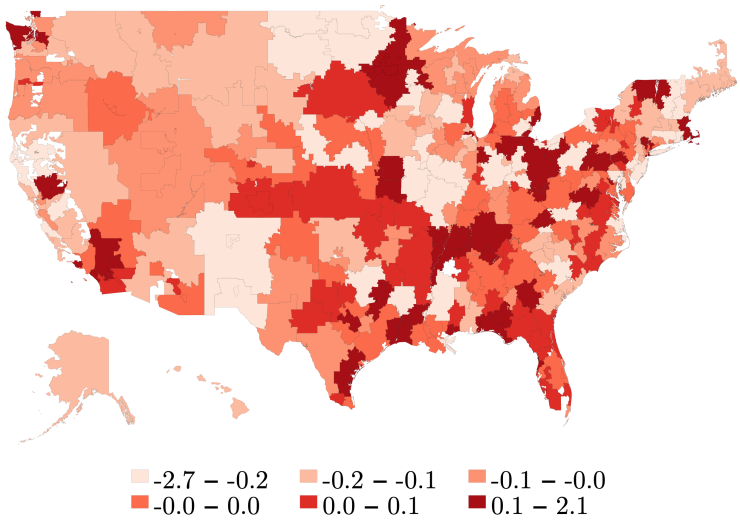

Notes: Plots the residual variation in the instrumental variable, conditional on the Lasso-selected covariates, across HRRs. The variable is standardized, so the units are in standard deviations.

Tables A5—A7 report the summary statistics and univariate regression coefficients for the 76 variables that form the basis of our potential control sets in the Lasso regressions.

Table A8 provides summary statistics and also illustrates several facts about covariate balance on a subset of variables that one might think ex ante could reflect important differences in preferences for branded statins across physicians, hospitals, and geographies.

Table A5: Covariate Summary Statistics

	mean	s.d.	meal β	mkt. share β
Physician Compare files				
Num. hospital affils.	3.6	1.4	0.07	–0.12
Num. secondary specialties	0.39	0.54	–0.03	–0.01
Num. zip-code affils.	1.6	1.3	–0.03	–0.02
At a group practice	0.97	0.18	0.07	–0.04
At a hospital	0.99	0.09	0.04	–0.02
At a AHA-matched hospital	0.99	0.12	0.05	–0.01
Participates in EHR	0.62	0.49	0.04	–0.02
Participates in e-Rx	0.70	0.46	0.12	–0.04
Is female	0.09	0.28	–0.09	0.02
Med. school grad. year	1,984	9.8	0.05	–0.07
Miles to nearest AMC	36.0	44.1	–0.03	–0.12
Participates in PQRS	0.52	0.50	0.00	–0.01
AMC & AMSA files				
Own AMSA $\times 1\{\text{Is AMC faculty}\}$	1.6	6.3	–0.25	0.13
Is AMC faculty	0.07	0.25	–0.24	0.13

Notes: Summary statistics are reported at the physician level. The "meal β" and "mkt. share β" columns report the coefficients from physician-level regressions of a dummy for meal receipt and the standardized market share of a drug, respectively, on the standardized variable.

The first column provides the mean of the variable in the sample of cardiologists who do not receive a meal. Below in brackets, it also provides the standard deviation of that variable across all cardiologists in our sample. These summary statistics help provide context for the magnitudes and variation for each variable, which helps among other things to interpret the balance results in the next two columns.

The next two columns present results on balance in the pseudo-experiment captured by our instrument by regressing the standardized variable on an indicator for whether the AMSA spillover instrument is above the median, and reporting the coefficient β that captures the relationship between a "high" AMSA spillover measure and a standard deviation change in the variable of interest (with standard errors below). The first column provides the coefficient from a univariate regression, and the second from a regression that also includes the controls selected by the Lasso in our main specification.

In the univariate regressions, almost all of the variables show that there are statistically significant differences between observations with low and high AMSA spillover instrument values, but very few show any economically meaningful differences. Of the 18 variables in the table, only three have differences greater than 0.2 standard deviations: cardiologist miles to nearest AMC, HSA Medicare Advantage enrollment, and HRR Medicare Advantage enrollment. One interpretation is that there are some baseline differences between low and

Table A6: Covariate Summary Statistics (cont'd)

	mean	s.d.	meal β	mkt. share β
Part B files				
Part B, allow amount, all services, avg.	78.7	34.8	0.13	0.03
Part B, allow amount, all services, *p10*	13.1	13.3	–0.08	0.05
Part B, allow amount, all services, *p25*	26.6	23.5	0.00	0.02
Part B, allow amount, all services, *p75*	100	35.1	0.13	0.07
Part B, allow amount, all services, *p90*	155	100	0.09	0.04
Part B, allow amount, cardio. med., avg.	55.0	34.9	0.10	0.05
Part B, allow amount, cardio. med., *p10*	13.1	14.6	–0.05	0.06
Part B, allow amount, cardio. med., *p25*	17.2	20.3	–0.04	0.05
Part B, allow amount, cardio. med., *p75*	77.2	64.4	0.09	0.05
Part B, allow amount, cardio. med., *p90*	131	82.8	0.15	0.04
Part B, allow amount, cardio. surg., avg.	92.5	357	0.07	–0.06
Part B, allow amount, cardio. surg., *p10*	57.4	217	0.06	–0.06
Part B, allow amount, cardio. surg., *p25*	63.0	231	0.06	–0.06
Part B, allow amount, cardio. surg., *p75*	114	549	0.06	–0.05
Part B, allow amount, cardio. surg., *p90*	148	754	0.06	–0.05
Part B, allow amount, emerg., avg.	95.1	19.2	0.04	0.11
Part B, allow amount, emerg., *p10*	61.4	24.0	–0.05	0.07
Part B, allow amount, emerg., *p25*	75.1	22.9	0.04	0.09
Part B, allow amount, emerg., *p75*	105	23.7	0.02	0.10
Part B, allow amount, emerg., *p90*	138	38.3	0.00	0.04
Part B, service count, all services	4,033	3,034	0.22	–0.13
Part B, service count, cardio. med.	1,783	1,365	0.14	–0.14
Part B, service count, cardio. surg.	78.4	237	0.09	0.03
Part B, service count, emerg.	1,627	1,128	0.28	–0.16
Part B, service count, new patients	63.0	56.5	0.15	–0.15
Part B, patient count, all services	956	588	0.29	–0.20
Part B, patient count, cardio. med.	1,455	1,046	0.15	–0.15

Notes: Summary statistics are reported at the physician level. The "meal β" and "mkt. share β" columns report the coefficients from physician-level regressions of a dummy for meal receipt and the standardized market share of a drug, respectively, on the standardized variable.

Table A7: Covariate Summary Statistics (cont'd)

	mean	s.d.	meal β	mkt. share β
Hospital level				
Admissions	14,065	12,587	–0.03	0.14
AMSA score, faculty-wgt.	1.1	4.6	–0.27	0.15
Cardiologists' market size, avg.	3,330	2,366	0.27	–0.18
Faculty share	0.04	0.13	–0.25	0.17
Beds	288	248	–0.03	0.14
Is teaching hospital	0.13	0.34	–0.19	0.19
Num. cardiologists	4.7	6.3	–0.10	0.19
Num. physicians	67.0	85.9	–0.13	0.16
HSA level				
Cardiologists' market size, avg.	3,663	2,178	0.23	–0.17
Faculty share	0.02	0.06	–0.19	0.17
Cardiac hospitalization rate	67.0	13.6	0.08	–0.11
Medicare Adv., eligibility	103,718	200,697	0.02	0.11
Medicare Adv., enrollment rate	21.6	13.6	–0.03	0.01
Medicaid, enrollment rate	22.6	8.3	–0.02	0.04
Uninsured rate	10.5	4.4	0.16	–0.10
Admissions, sum	147,235	157,606	–0.08	0.20
Beds, sum	3,060	3,209	–0.07	0.19
Cardiologists, sum	11.4	22.6	–0.09	0.16
Teaching hospitals, sum	1.6	2.5	–0.11	0.22
Physicians, sum	226	400	–0.10	0.15
HRR level				
Cardiologists' market size, avg.	3,146	1,366	0.20	–0.17
Cardiac hospitalization rate, avg.	66.1	12.2	0.14	–0.08
Faculty share, avg.	0.06	0.07	–0.16	0.13
Medicare Adv., eligibility, avg.	104,213	199,390	0.02	0.11
Medicare Adv., enrollment rate, avg.	20.3	11.5	–0.03	0.04
Medicaid, enrollment rate, avg.	22.6	6.8	0.04	–0.01
Uninsured rate, avg.	10.6	4.3	0.18	–0.11
Admissions, sum	2,319,274	2,853,892	–0.08	0.11
Beds, sum	47,300	55,813	–0.06	0.11
Cardiologists, sum	130	131	–0.07	0.15
Teaching hospitals, sum	26.9	38.9	–0.14	0.15
Physicians, sum	2,607	2,515	–0.10	0.12
Advertising: Zipcode-drug level				
Ad volume, duration	63,336	23,194	–0.65	0.27
Ad volume, spending	34,428	49,605	–0.17	0.25
Ad volume, units	1,057	387	–0.66	0.27

Notes: Summary statistics are reported at the level of variation in the variable. Rates are reported on a 0–100 (p.p.) scale. All HRR-level variables reported as "avg." are averages of the HSAs within the HRR due to the level of data-reporting. The "meal β" and "mkt. share β" columns report the coefficients from physician-level regressions of a dummy for meal receipt and the standardized market share of a drug, respectively, on the standardized variable.

high AMSA regions in the size and density of cardiology activity. This does not achieve the precise level of balance we would see in, say, a randomized control trial, but our exclusion restriction requirement is conditional on controls. We find it particularly reassuring that the differences between cardiologists' patients (captured in Table A8 as "Part B allow amount" variables) between high and low AMSA spillover observations are economically small. These are also variables that one might think would be independently correlated with pharmaceutical sales force allocation, but they are also ones that we think reasonable controls should be able to capture.

Once we add our preferred controls in the next column, many of the differences between the low and high AMSA regions shrink. In particular, now only a single variable has a difference greater than 0.2 standard deviations: cardiologist miles to nearest AMC is now about 0.3 standard deviations (or 15 miles, at a mean of 49 miles) different between the two samples.

In the end, our instrument shows quite good balance across a large number of potentially important variables at the individual cardiologist, hospital, and regional levels. This balance improves with our preferred controls. Indeed, many of the variables here (in particular the cardiologist miles to nearest AMC) are themselves selected as controls, as illustrated by the final column. Thus we are reassured that our chosen instrument is largely exogenous with respect to other potentially important determinants of payments and prescribing, particularly after including controls.

G.2 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.

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 A8: Instrument Balance across Select Covariates

	mean if no meal [s.d., full samp]	AMSA β	AMSA β + controls	share of subsamples selected in
Cardiologist: At a AHA-matched hospital	0.9921	–0.1018	0.0027	0.01
	[0.1106]	(0.0021)	(0.0088)	
Cardiologist: Market size	7.734	–0.1537	0.0073	1.00
	[0.7984]	(0.0019)	(0.0076)	
Cardiologist: Med. school grad. year	1984	–0.0597	0.0417	0.94
	[9.672]	(0.0023)	(0.0104)	
Cardiologist: Miles to nearest AMC	48.78	–0.7209	–0.2976	0.98
	[43.87]	(0.0012)	(0.0203)	
Cardiologist: Part B, allow amount, all services, *p75*	98.32	0.1812	0.0157	0.00
	[34.25]	(0.002)	(0.0065)	
Cardiologist: Part B, allow amount, all services, mean	77.98	0.0691	–0.0061	0.00
	[34.52]	(0.0018)	(0.0071)	
Cardiologist: Part B, allow amount, cardio. med., *p75*	76.02	0.0726	–0.0207	0.00
	[64.14]	(0.002)	(0.0091)	
Cardiologist: Part B, allow amount, cardio. med., mean	54.55	0.0542	–0.0275	0.00
	[34.34]	(0.002)	(0.0094)	
Cardiologist: Participates in EHR	0.6502	–0.1287	0.023	0.00
	[0.484]	(0.0022)	(0.0084)	
Cardiologist: Participates in e-Rx	0.7411	–0.133	–0.1114	0.20
	[0.4503]	(0.002)	(0.0047)	
Hospital: Beds	437.3	–0.0838	–0.1112	0.00
	[328.3]	(0.0022)	(0.0143)	
Hospital: Is teaching hospital	0.2124	0.0937	–0.1202	1.00
	[0.4194]	(0.002)	(0.0112)	
HSA: Cardiac hospitalization rate	65.85	–0.039	0.1267	1.00
	[11.38]	(0.002)	(0.0102)	
HSA: Medicaid, enrollment	21.85	0.1986	–0.0294	0.18
	[8.093]	(0.0022)	(0.0087)	
HSA: Medicare Adv., enrollment	23.43	0.2109	0.1023	1.00
	[13.27]	(0.0021)	(0.0407)	
HRR: Cardiac hospitalization rate	67.56	–0.1639	–0.1751	1.00
	[11.62]	(0.0019)	(0.0098)	
HRR: Medicaid, enrollment	22.72	0.0249	0.0177	0.97
	[6.761]	(0.0022)	(0.021)	
HRR: Medicare Adv., enrollment	19.05	0.3698	0.0573	0.00
	[11.21]	(0.0019)	(0.0086)	

Notes: The "AMSA β" reports the coefficient from a regression of the standardized variable on an indicator that the AMSA I.V. is above the median value, without (second column) or with controls (third column).

Table A9: Alternative Demand Specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
θ^p	–0.00816	–0.00776	–0.00032	–0.00816	–0.00835	–0.01247	–0.00805	–0.00708	–0.00839
	(0.00018)	(0.00014)	(0.00003)	(0.00005)	(0.00018)	(0.00010)	(0.00018)	(0.00017)	(0.00019)
λ_{statins}	0.339	0.157	0.961	0.339	0.334		0.364	0.430	0.322
	(0.0129)	(0.0135)	(0.0009)	(0.0039)	(0.0132)		(0.0127)	(0.0127)	(0.0137)
$\lambda_{\text{strong statins}}$							–0.038		
							(0.0084)		
mean(η^p)	–0.209	–0.160	–0.125	–0.209	–0.212	–0.222	–0.208	–0.207	–0.210
s.d.(η^p)	0.222	0.168	0.141	0.222	0.226	0.230	0.221	0.222	0.223
F stat.	390.6	525.1	12,451.5	390.6	336.2	36,112.6	282.3	363.6	372.1
$R^2(\delta_{djt} : \psi_{dj} - \theta^p p)$	0.84	0.86	0.20	0.84	0.84	0.88	0.83	0.81	0.85
mean($\psi_{dj}/\lvert\theta^p\rvert$)$_{\text{str. stat.}}$	–306.7	–363.3	–5,335.7	–306.7	–302.5	–236.1	–308.4	–316.2	–298.4
mean($\psi_{dj}/\lvert\theta^p\rvert$)$_{\text{other gen.}}$	–327.1	–382.4	–5,439.0	–327.1	–322.7	–255.3	–326.8	–336.8	–320.3
s.d.($\psi_{dj}/\lvert\theta^p\rvert$)	80.4	98.5	1,407.3	80.4	78.7	69.6	79.9	79.3	85.7
Specification									
FE level	*dj*	*dj*	*dj*	*dj*	*dj*	*dj*	*dj*	*dj*	*dj*
price instrument	Y		Y	Y	Y	Y	Y	Y	Y
nest instrument	Y	Y		Y	Y	Y	Y	Y	Y
q perturb.	Y	Y	Y		Y	Y	Y	Y	Y
drop faculty					Y				
alt. market size def'n.								A1	A2
N obs.	100,763	100,763	100,763	100,763	94,623	100,763	100,763	100,435	100,041

Notes: Replicates the price and nest regression using the preferred specification (Col. 1) and seven alternate specifications. Parameter estimates based on Eq. 8. Standard errors for the main parameters (θ^p and λ), in parentheses, are based on the standard deviation of the point estimates from the perturbed-bootstrap samples.

G.3 Alternative MTE Results and Specifications

The robustness of our main results to the selection of controls is shown in Figure A4. The top Panel shows the first stage coefficient on our AMSA spillover IV and the average treatment effects implied by the MTE estimates from a range of alternative specifications. The bottom Panel plots the distributions of MTEs for the same specifications. Specification (1) is our preferred specification where we employ the Lasso to select from among the full set of possible controls. Specification (2) excludes cardiologist-level controls (including the "Part B" controls for the size and intensity of patient panels) from the set that may be selected by the Lasso. Specification (3) excludes hospital level controls. Specification (4) excludes advertising controls. Specification (5) excludes all market (HSA and HRR) level controls. Specification (6) omits the Lasso selection step and includes all controls. Specification (7) includes the Lasso selection step, but allows selection from among a larger set of 283 possible controls: the full set of controls, plus log and quadratic transformations thereof. In specifications (8) and (9), we matched every physician with an above-median AMSA spillover instrument to another physician with a below-median AMSA instrument. Specification (8) matches physicians only on distance to the nearest AMC (M1). Specification (9) matches physicians on all the variables shown in Table A8 (M2). In each case, we used propensity score matching with replacement and imposed the requirement that every matched pair be within 0.2 standard deviations on all match variables.

With the exception of specification (5), all alternative specifications are qualitatively and quantitatively similar, and the MTE curves are essentially contained within the 95% confidence interval for our preferred specification (1). Specification (5) has a negative point estimate and positively sloped MTE curve. Examining the summary statistics for these covariates and their relationships with meal receipt and prescribing market share in Table A7, we observe that a number of market-level (i.e., HSA- or HRR-level) covariates have opposite-signed relationships with meal receipt and prescribing market share. For example, cardiologists in HSAs and HRRs where other cardiologists have larger average market sizes are more likely to receive meals, but have lower branded statin market shares. Thus, we consider it important either to allow the Lasso to select and control for market level covariates that are correlated with both promotional activities and prescribing patterns, or to match physicians with high and low AMSA IVs in similar markets. But the exact specifics as to how we control for such correlates do not influence the results, and beyond these correlates, our estimates are largely insensitive to controls for cardiologist characteristics, cardiologist patient panels, hospital characteristics, and advertising activity.

Figure A4: Sensitivity of MTE Curves to Covariates

(a) Treatment Effect Point Estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
				First stage estimate					
AMSA I.V.	–0.030	–0.032	–0.029	–0.034	–0.026	–0.022	–0.031	–0.034	–0.029
	(0.0011)	(0.0011)	(0.0011)	(0.0012)	(0.0005)	(0.0011)	(0.0016)	(0.0012)	(0.0012)
				MTE-based estimates					
ATE	0.221	0.291	0.296	0.178	–0.085	0.098	0.370	0.221	0.154
	(0.0940)	(0.0928)	(0.0902)	(0.0774)	(0.0531)	(0.1072)	(0.1040)	(0.0785)	(0.1110)
ATT	0.344	0.428	0.393	0.294	–0.109	0.322	0.509	0.472	0.482
	(0.0992)	(0.0995)	(0.0997)	(0.0864)	(0.0584)	(0.1030)	(0.1128)	(0.0847)	(0.1199)
ATUT	0.075	0.140	0.188	0.037	–0.068	–0.147	0.206	–0.060	–0.183
	(0.1285)	(0.1182)	(0.1277)	(0.1073)	(0.0663)	(0.1613)	(0.1391)	(0.1141)	(0.1460)
F stat.	761.2	808.6	710.8	856.0	2,547.5	391.3	381.0	761.1	558.9
N obs.	21,127	21,127	21,127	21,127	21,127	21,127	21,127	16,172	14,776
X incl.	Y	Y	Y	Y	Y	Y	Y	Y	Y
NPI-level	Y		Y	Y	Y	Y	Y	Y	Y
hosp.-level	Y	Y		Y	Y	Y	Y	Y	Y
advert.	Y	Y	Y		Y	Y	Y	Y	Y
HSA- & HRR-level	Y	Y	Y	Y		Y	Y	Y	Y
Lasso	Y	Y	Y	Y	Y		Y	Y	Y
Non-lin. X							Y		
Matched sub-samp.								M1	M2

(b) Treatment Effect Distributions

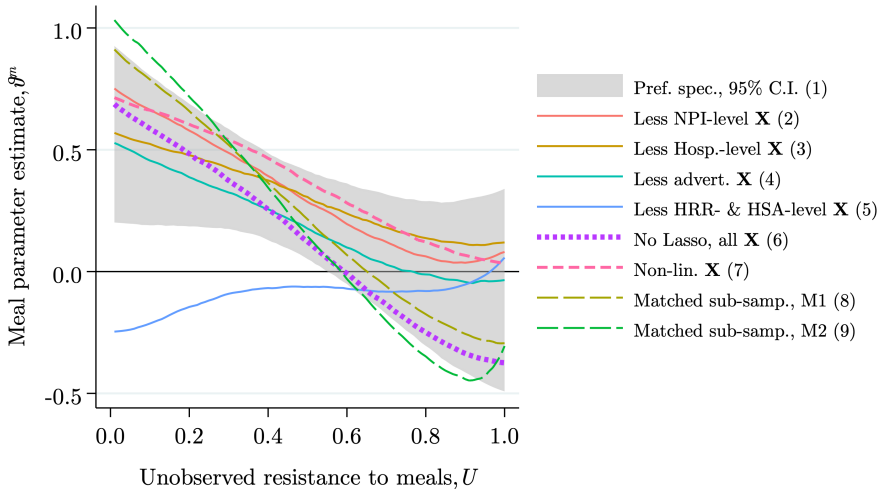

Notes: The columns in the table correspond to the MTE curves indicated by the legend in the figure; see the accompanying text for details of the alternative specifications. The shaded gray indicates 95% confidence intervals.

G.4 Exploration of Treatment Effects of Payments

In the following analyses, we examine the robustness of our results to sample restrictions, market definition, and model implementation, and we investigate whether cardiologists seem to respond to payments beyond meals or to higher values of meal payments.

Figure A5 recreates the MTE curves and displays the ATE/ATT/ATUT estimates corresponding to our preferred specification shown in the main text as column (1), as well as nine alternative specifications (2–10). Specification (2) does not perturb the annual claim quantities at the use-case level (the preferred model does). Specification (3) excludes all AMC faculty from the entire estimation routine (the preferred model includes them). Specification (4) uses the alternative market definition of all cardiovascular drugs (A1). Specification (5) excludes other large cardiovascular classes with branded drugs from the market size calculation (A2).[60] Specification (6) replaces our preferred nonparametric local linear specification of the control functions in the MTE model with a second degree (quadratic) local polynomial (Andresen, 2018). Each of these modeling alternatives results in similar first stage and average treatment effect estimates in Panel (a), and their corresponding MTE curves are contained within the 95% confidence interval of the preferred estimate in Panel (b).

Figure A5, columns (7)-(10) in Panel (a), and the MTE curves in Panel (c), explore the nature of the payment response beyond what might be captured in our focal meal dummy. Column (7) drops physicians in the top 5% of the meal dollars distribution (above $203); column (8) drops the top 10% (above $147); column (9) drops the top 25% (above $85). Thus, these specifications estimate the effect of receiving meals of *successively smaller dollar values*. Column (10) replaces our preferred indicator for meal-based relationships with an indicator for receiving any kind of payment (e.g., meals, consulting, speaking, travel, or research). The MTE curves and implied average treatment effects in columns (7)-(10) are all quite similar to our full sample estimates. The results are slightly smaller for "Any General Payment," but are not statistically significantly different from our main specification. In sum, we do not find evidence that physicians are more responsive to higher payments, or that their response to payments depends on the qualitative nature of payments.

[60] A potential endogeneity problem could arise if our instrument shuts down meals for other, non-statin cardiovascular drugs, such that our instrument would have a direct unobserved negative effect on market size and, in turn, a positive effect on $ln(s_j/s_0)$. This would bias our estimated meal coefficient toward zero. To investigate this possibility, we eliminate non-statin cardiovascular drug classes containing drugs with the highest potential for payments. There were five drugs with at least 2% shares of the overall cardiovascular class by sales and quantity: Cardizem, Coreg, Diovan, Dutoprol, and Exforge, contained in non-statin ATC4 categories C05A (vasoprotectives), C07A (beta blockers), and C09C and C09D (two types of angiotensin II receptor blockers). Collectively, all drugs in these ATC4 categories accounted for 41% of cardiovascular prescriptions for the average sample physician; for perspective, Crestor and Lipitor together accounted for 19% of cardiovascular prescriptions.

Figure A5: Additional Alternative MTE Specifications

(a) Treatment Effect Point Estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
				First stage estimate						
AMSA I.V.	–0.030	–0.030	–0.037	–0.029	–0.029	–0.030	–0.029	–0.028	–0.026	–0.030
	(0.0011)	(0.0011)	(0.0012)	(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0011)
				MTE-based estimates						
ATE	0.221	0.229	0.322	0.164	0.235	0.262	0.252	0.290	0.291	0.187
	(0.0940)	(0.0926)	(0.0888)	(0.0804)	(0.1125)	(0.0966)	(0.0926)	(0.1076)	(0.1233)	(0.0906)
ATT	0.344	0.353	0.481	0.271	0.338	0.336	0.335	0.368	0.375	0.307
	(0.0992)	(0.0960)	(0.1037)	(0.0854)	(0.1215)	(0.1041)	(0.1000)	(0.1109)	(0.1280)	(0.0999)
ATUT	0.075	0.077	0.168	0.050	0.124	0.173	0.154	0.234	0.214	0.054
	(0.1285)	(0.1267)	(0.1076)	(0.1123)	(0.1534)	(0.1309)	(0.1352)	(0.1535)	(0.1823)	(0.1250)
F stat.	761.2	742.0	973.2	740.2	647.1	760.9	709.1	659.5	541.2	757.0
N obs.	21,127	21,127	19,785	21,069	20,964	21,127	20,576	20,023	18,364	21,127
No q perturb.		Y								
Drop faculty			Y							
Alt. mkt. size				A1	A2					
2-deg. loc. poly.						Y				
Drop $p5$%-$							Y			
Drop $p10$%-$								Y		
Drop $p25$%-$									Y	
Any type										Y

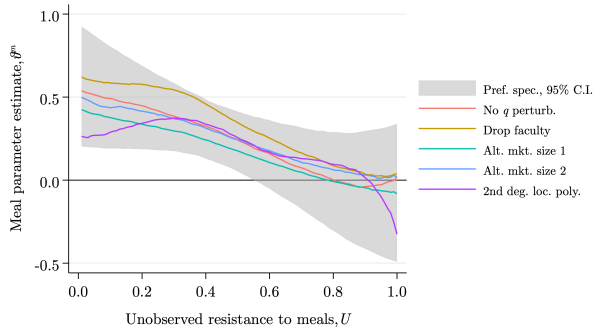

(b) Treatment Effect Distributions

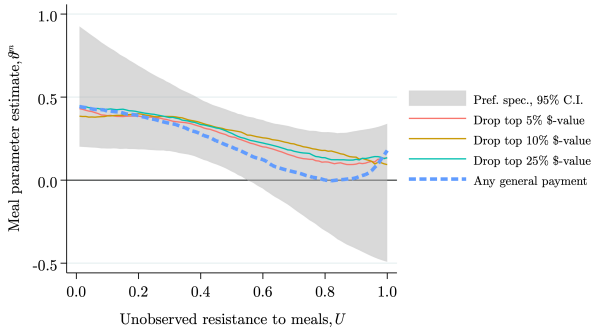

(c) Treatment Effect Distributions

Notes: The columns in the table correspond to the MTE curves indicated by the legends in the figure; see the accompanying text for details of the alternative specifications. The shaded gray indicates 95% confidence intervals.

H Additional Counterfactual Results and Robustness

Table A10: Equilibrium Quantity and Price Effects of Meal Payments (2012)

	Observed	Ban, fix p	Ban	Ban, $p = mc$	No Ban, $p = mc$
$Q_{statins}$	0.166	0.162	0.162	0.174	0.179
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
$Q_{Lipitor}$	0.052	0.053	0.053	0.061	0.060
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
$Q_{Crestor}$	0.022	0.016	0.016	0.023	0.032
	(0.000)	(0.001)	(0.001)	(0.002)	(0.000)
$OOP_{statins}$	17.53	17.53	17.41	2.15	2.15
	(0.26)	(0.26)	(0.25)	(0.1)	(0.1)
$POS_{statins}$	62.39	62.39	61.93	85.83	86.76
	(0.64)	(0.64)	(0.62)	(2.23)	(2.19)

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. $N = 105,865$ doctor-drug-brand-year observations with standard errors clustered at the doctor level ($N_d = 11,870$) via delete-220 jackknife and state level via delete-7 jackknife.

Table A11: Welfare and Counterfactual Estimates – Supplement to Figure 6

	ε^{de}	-350	-100	-50	0
Total Surplus	Observed, 2011	76.45	37.21	29.37	21.52
		(0.66)	(0.51)	(0.48)	(0.46)
	Ban, 2011	-2.38	-0.53	-0.15	0.20
		(0.48)	(0.08)	(0.03)	(0.04)
	Observed, 2012	4.64	2.48	2.05	1.62
		(0.19)	(0.18)	(0.17)	(0.17)
	Ban, 2012	3.56	2.28	2.02	1.76
		(0.12)	(0.13)	(0.16)	(0.19)
Consumer Surplus	Observed, 2011	74.54	35.32	27.48	19.63
		(0.64)	(0.49)	(0.47)	(0.46)
	Ban, 2011	-2.28	-0.41	-0.06	0.30
		(0.40)	(0.03)	(0.04)	(0.13)
	Observed, 2012	5.01	2.84	2.41	1.97
		(0.17)	(0.17)	(0.17)	(0.17)
	Ban, 2012	3.94	2.64	2.38	2.13
		(0.09)	(0.17)	(0.20)	(0.24)
Consumer Surplus (-Transfers)	Observed, 2011	67.16	27.93	20.08	12.22
		(0.63)	(0.49)	(0.48)	(0.49)
	Ban, 2011	-0.95	0.86	1.23	1.59
		(0.13)	(0.24)	(0.32)	(0.40)
	Observed, 2012	6.87	4.72	4.28	3.85
		(0.18)	(0.18)	(0.18)	(0.18)
	Ban, 2012	6.68	5.36	5.11	4.86
		(0.20)	(0.35)	(0.39)	(0.43)
Producer Surplus	Observed, 2011	9.27	9.27	9.27	9.27
		(0.11)	(0.11)	(0.11)	(0.11)
	Ban, 2011	-1.40	-1.40	-1.40	-1.40
		(0.36)	(0.36)	(0.36)	(0.36)
	Observed, 2012	-2.22	-2.22	-2.22	-2.22
		(0.04)	(0.04)	(0.04)	(0.04)
	Ban, 2012	-3.11	-3.11	-3.11	-3.11
		(0.23)	(0.23)	(0.23)	(0.23)

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. $N = 105,865$ doctor-drug-brand-year observations with standard errors clustered at the doctor level ($N_d = 11,870$) via delete-220 jackknife and state level via delete-7 jackknife.

Table A12: Robustness of Welfare Estimates to Modeling Assumptions

	ε^{de}	-350	-100	-50	0
Consumer Surplus, Ban 2011 (Retail)	Baseline	-2.28	-0.41	-0.06	0.30
		(0.40)	(0.03)	(0.04)	(0.13)
	Rebates	-2.25	-0.40	-0.06	0.30
		(0.39)	(0.03)	(0.04)	(0.13)
	MC	-2.25	-0.40	-0.06	0.31
		(0.39)	(0.03)	(0.04)	(0.13)
	Pricing	-2.57	-0.70	-0.31	0.03
		(0.46)	(0.06)	(0.02)	(0.06)
	ϵ_{DE}	-3.53	-0.79	-0.25	0.30
		(2.93)	(0.92)	(0.46)	(0.13)
	High θ_p	-2.21	-0.51	-0.15	0.18
		(0.41)	(0.06)	(0.01)	(0.08)
	Low θ_p	-1.54	0.30	0.68	1.06
		(0.10)	(0.31)	(0.41)	(0.49)
Consumer Surplus, Ban 2011 (-POS Transfers)	Baseline	-0.95	0.86	1.23	1.59
		(0.13)	(0.24)	(0.32)	(0.40)
	Rebates	-0.92	0.88	1.24	1.60
		(0.11)	(0.25)	(0.33)	(0.41)
	MC	-0.90	0.89	1.25	1.62
		(0.12)	(0.25)	(0.33)	(0.41)
	Pricing	-1.44	0.24	0.51	0.73
		(0.22)	(0.11)	(0.17)	(0.21)
	ϵ_{DE}	-2.22	0.46	1.04	1.59
		(2.92)	(0.88)	(0.49)	(0.40)
	High θ_p	-0.92	0.79	1.13	1.48
		(0.16)	(0.20)	(0.27)	(0.34)
	Low θ_p	-0.41	1.50	1.88	2.26
		(0.12)	(0.59)	(0.66)	(0.74)
Producer Surplus, Ban 2011	Baseline	-1.40	-1.40	-1.40	-1.40
		(0.36)	(0.36)	(0.36)	(0.36)
	Rebates	-1.41	-1.41	-1.41	-1.41
		(0.36)	(0.36)	(0.36)	(0.36)
	MC	-1.44	-1.44	-1.44	-1.44
		(0.37)	(0.37)	(0.37)	(0.37)
	Pricing	-1.16	-0.94	-0.83	-0.64
		(0.31)	(0.25)	(0.22)	(0.18)
	ϵ_{DE}	-1.40	-1.40	-1.40	-1.40
		(0.36)	(0.36)	(0.36)	(0.36)
	High θ_p	-1.40	-1.40	-1.40	-1.40
		(0.35)	(0.35)	(0.35)	(0.35)
	Low θ_p	-1.27	-1.27	-1.27	-1.27
		(0.32)	(0.32)	(0.32)	(0.32)

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 A11), and alternative specifications: "Rebates" (alternative rebates as described in Appendix E); "Marginal Costs" (extreme alternative assumption that $mc = 0$); "Pricing" ($\alpha^{de} = 1$ in model in Appendix C.5 with insurer sophistication, whereas baseline specification sets $\alpha^{de} = 0$); and "ϵ_{DE}" (an alternative specification with $\varepsilon_d^{de} = \gamma^{de} * \bar{\theta}_d^m$ rather than fixed ε^{de} across all physicians). For the $\varepsilon_d^{de} = \gamma^{de} * \bar{\theta}_d^m$ specifications, the column value of ε^{de} is the average across sample physicians, given their average meal responsiveness $\bar{\theta}_d^m$. High (Low) θ_p changes the price sensitivity parameter in the utility model so that the elasticities are shifted one SD higher (lower) than our estimated elasticities (SD based on drug elasticity estimates in (Einav et al., 2018)). $N = 105,865$ doctor-drug-brand-year observations with standard errors clustered at the doctor level ($N_d = 11,870$) via delete-220 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	0.16	0.16	0.16	0.16	0.16	0.16	0.16	0.16
		(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	0.03	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)
Q Crestor	Observed	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02
		(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	18.73	18.73	18.73	18.73	18.67	18.45	18.36	18.21
		(0.20)	(0.20)	(0.20)	(0.20)	(0.23)	(0.25)	(0.26)	(0.26)
	Ban	-0.20	-0.20	-0.20	-0.20	0.70	1.59	2.07	2.92
		(0.03)	(0.03)	(0.03)	(0.03)	(0.21)	(0.48)	(0.63)	(0.92)
POS Statins	Observed	74.91	74.91	74.91	74.91	74.60	73.69	73.30	72.68
		(0.74)	(0.74)	(0.74)	(0.74)	(0.79)	(0.94)	(1.05)	(1.15)
	Ban	-0.84	-0.84	-0.84	-0.84	2.95	6.65	8.67	12.29
		(0.13)	(0.13)	(0.13)	(0.13)	(0.86)	(1.94)	(2.54)	(3.75)

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 ε^{de}. $N = 105,865$ doctor-drug-brand-year observations with standard errors clustered at the doctor level ($N_d = 11,870$) via delete-220 jackknife and state level via delete-7 jackknife.