

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

Matthew Grennan

Emory University, UC-Berkeley, Haas School & NBER, USA

Kyle R. Myers

Harvard Business School & NBER, USA

Ashley Swanson

UW-Madison & NBER, USA

and

Aaron Chatterji

Duke University, Fuqua School & NBER, USA

First version received March 2021; Editorial decision April 2024; Accepted September 2024 (Eds.)

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 AMC CoI 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, whether consumers are harmed by payments depends on whether there is substantial under-prescribing due to behavioural or other frictions. In a final exercise, we calibrate such frictions using clinical data and estimate that payments benefit consumers in this case study.

Key words: Expert agents, Conflict of interest, Health care, Welfare, Decision errors

JEL codes: L00, I11, I18

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 (David *et al.*, 2010; Anagol *et al.*, 2017; Egan *et al.*, 2019; Bhattacharya *et al.*, 2020; 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 behaviour 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 behavioural biases) that might drive a wedge between the treatment a physician chooses and the treatment that maximises 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 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 generalisability 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% of the prescriptions in our data were written by physicians who received payments from at least one firm. Lipitor and Crestor made up nearly 40% 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 (OOP) price the patient must pay, the effect of any payment interaction

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), and 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.

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

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 payment 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 AMC's 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. 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.⁵ 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.⁶ We use our AMSA instrument to estimate the distribution of marginal treatment effects (MTEs; Heckman and Vytlacil, 2007) of payments

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 the prescribing of promoted drugs at the institutions that impose them.

5. As discussed in further detail in Section 2.4 and Appendix G.3, we do find it particularly important for the purposes of this cross-sectional identification strategy to compare physicians practicing in observably similar markets.

6. 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 (Shapiro, 2018a; Agha and Zeltzer, 2019; Carey *et al.*, 2020). 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).

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).

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% for the average physician, which is roughly equivalent to the impact of a \$28 decline in OOP price, or 0.37 standard deviations in the prescribing heterogeneity across physicians. However, the effect is 0.71 standard deviations for a physician in the 90th percentile of the treatment effect distribution, while the effect in the 10th 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 analyse how prices are determined in equilibrium.⁷ 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 behavioural biases (Baicker *et al.*, 2015). In our model, payments could reinforce or counteract such frictions.⁸ 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% 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 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

7. Prescription drug utilisation 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 versus 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 enrollee premiums.

8. 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).

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 behavioural 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 IV 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 (Levitt and Syverson, 2008; Schneider, 2012; Anagol *et al.*, 2017; Egan *et al.*, 2019; Bhattacharya *et al.*, 2020) and in particular to the literature on drivers of physician treatment recommendations (Gruber and Owings, 1996; Iizuka, 2012; Clemens and Gottlieb, 2014; Ho and Pakes, 2014; Dickstein, 2017). 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; Sinkinson and Starc, 2019). Our equilibrium model contributes to an evolving literature on the price effects of factors that shift drug utilisation (Dafny *et al.*, 2024; Tuncel, 2024). 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 U.S., and in particular, the Medicare statin market in 2011–12. 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, 2011–12

With prescription drugs accounting for more than 15% of personal health care expenditures, and with 72% 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% and 27% 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% of cardiologist-prescribed revenue for both Pfizer and AstraZeneca, making it likely that they are an important subject of any firm interactions with cardiologists.⁹ Also, although cardiologists accounted for only about 10% 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% (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 (Goldman *et al.*, 2007; Chandra *et al.*, 2010; Abaluck *et al.*, 2018). 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

9. In contrast, in interactions with primary care providers (PCPs) in 2011–12, Pfizer might have promoted Celebrex, Enbrel, Lipitor, Lyrica, Norvasc, Prevnar, Pristiq, Viagra, and Zyxos, and AstraZeneca might have promoted Synagis, Toprol, Seroquel, Atacand, Nexium, Prilosec, and Symbicort. Consistent with this, after Lipitor went off patent, payments from Pfizer to cardiologists sharply declined, but no such trend break was observed for PCPs. Results for all of our main analyses run on PCPs are available by request.

(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). Because 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 degree of confidence in physicians' practice primary practice locations.¹⁰

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 allow 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 analysing days supplied as the unit of quantity (Starc and Swanson, 2021).

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 nonstatin 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.¹¹

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 fifth 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 nonzero observations (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% of Part D statin prescriptions and expenditures during 2011–12. Appendix [Table A2](#) details the impact of these sample restrictions on key summary statistics.

Table 1 summarises 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%. The effect of entry by generic atorvastatin in December 2011 was

10. 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 2005–10; 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.

11. In order to avoid any mechanical effect of statin prescribing on market size, we proxy for market size using nonstatin 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 nonstatin 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		OOP 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. OOP 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.

dramatic. In its first full year of availability, this new alternative accounted for more than 25% of cardiologists' statin claims.

The remaining columns of Table 1 summarise prices. In 2011, Lipitor and Crestor 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. In the U.S., it is common practice for insurers and pharmacy benefit managers (PBMs) to negotiate confidential rebates with manufacturers, typically as a discount off the POS price, which accrues to the insurer/PBM. As in most studies of pharmaceuticals, it is impossible for us to observe negotiated rebates or 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.* [Sood et al., 2017](#); [Yu et al., 2018](#); [Kakani et al., 2020](#); see Appendix E for details), suggest that 55–68% 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.¹²

12. 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.

2.3. *Payments to physicians*

More than 85% 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).¹³ These documents are the basis of our payments data, which were generously shared by Kyruus, Inc.¹⁴

Table 2 summarises 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.¹⁵ 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, given that they are clearly the dominant form of payment in this setting.

Sixty-seven percent of physicians, representing 77% of cardiovascular prescriptions in our sample, received a meal from at least one of the branded statin manufacturers. Meal-based relationships are highly persistent over time: for the firm-years in our estimation sample, 73% of physicians receiving a meal in year t also receive a meal in year $t + 1$. Further, there is not a large amount of variation in the dollar amount of meals when outliers are excluded: the 90th percentile of the distribution of meal dollar values across (nonzero) observations at the physician-firm-year level was less than \$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. The cross-sectional indicator for a meal also 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

13. 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.

14. 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.

15. Among physician-firm pairs that involved meals in our sample, 17% also involved nonmeal payments, but among those pairs that did not involve meals, less than one-half of a percent involved nonmeal payments.

TABLE 2
Payments summary statistics

	All types		Meals		Travel, speak, or consult		Research	
	Raw	Weighted	Raw	Weighted	Raw	Weighted	Raw	Weighted
Panel (a): Fraction of cardiologists receiving payments, by type								
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.

we instead use an indicator for receipt of *any* type of payment (*e.g.* meals, consulting, speaking, travel, or research).

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). Firms’ marketing models can be very detailed and data-driven, and pharmaceutical sales forces maintain rich databases on prescribers’ practice characteristics, prescribing behaviour, 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.¹⁶ Costs include the labour 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: AMCs’ restrictive access policies, specialty-specific physician employment by hospitals and health systems that have central purchasing or otherwise limit physicians’ autonomy, pressures on physicians that limit available time for firm interaction, etc. (Khedkar and Sturgis, 2015). As a result, the optimal allocation of sales force resources to regions and to physicians within those regions, taking into account economies of scale and the factors that determine the expected costs and benefits of

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




FIGURE 1

Regional variation in prescribing and meal payments, 2011 (a) Market Share (p.p.). (b) Share Doctors with Meal (p.p.), and (c) Market Share (p.p.), by Meal Status

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.

interaction (some of which are unobservable to the researcher), is a complex problem into which firms invest significant data, analytic, and personnel resources.

2.4. Regional variation and conflict-of-interest policies

Figure 1 documents the geographic variation in utilisation and meal payments across the U.S. Aggregating to the HRR level, Panel (a) plots the utilisation 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 90th percentile HRR per prescribing used roughly twice as many branded strong statins as those in the 10th percentile. Likewise, cardiologists in the 90th percentile HRR per payments were roughly three times as likely to receive a meal compared to those in the 10th 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 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, AMCs across the U.S. had a wide range of policies intended to prevent conflicts of interest by limiting

physician-firm relationships. We hypothesise 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 organisation studies (Hastings *et al.*, 2017), and to a broader literature that examines behaviour 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. As noted above, these analyses will typically include a number of variables that impact the expected benefits and costs of accessing physicians. In many regions of the country, local AMCs employ a significant fraction of physicians (particularly specialists like the cardiologists we focus on), and thus the cost or feasibility of accessing physicians associated with the AMC will be an important determinant of sales force allocation to the region.

We obtain measures of AMCs' CoI policies from the AMSA CoI scorecard. We link physicians to AMCs using the Association of American Medical Colleges (AAMC) faculty roster.¹⁷ 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.¹⁸ We hypothesise that regions where AMCs have strong CoI 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 hypothesise 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 nonfaculty 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'

17. 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.

18. 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 noncompliance. For each category except Oversight and NonCompliance, 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.




FIGURE 2

AMSA-scored CoI policies and meals (a) Across Faculty. (b) Across Hospitals and (c) Across HRRs

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. Based on 1,395, 2,775, and 21,642 physician-drug observations for Panels (a-c), respectively. The faculty weight is the share of all doctors in the hospital or region that are faculty.

behaviours. 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.¹⁹

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.²⁰

2.4.2. Physician-level and regional characteristics. The primary concern with using CoI policies as IV is that they vary only in the cross-section, and the exclusion restriction may fail due to direct effects of CoI policies on norms regarding prescribing, or due to unobservable factors correlated with selection into more restrictive policies (see discussion in [Larkin *et al.*, \(2017\)](#)).²¹ To help ensure that our identifying variation is driven by spillovers from CoI policies rather than these other factors (*e.g.* preferences, market structure, etc.), 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 organisations and practice

19. Personal communications: George Chressanthos 5 January 2018 and Pratap Khedkar 15 February 2018.

20. The geographic distribution of the instrumental variable is shown in Appendix [Figure A3](#).

21. 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.

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 10th/25th/75th/90th 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 hospitalisation rates from the Behavioural 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 utilisation 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 discussed in detail in Appendix G.3, our estimated treatment effects are also qualitatively and quantitatively robust as we vary the set of included controls, provided that we control for regional variables. That is, our estimates are stable conditional on using market-level variables to ensure that we are comparing physicians practicing in similar markets. Our results are also similar when we 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 POS 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 OOP 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 $j \in \mathcal{J} = \{0, 1, \dots, J\}$, for use case i (a doctor/patient/visit combination) in each market defined by doctor d in year t , be given by: $u_{idjt} = \delta_{djt} + \varepsilon_{idjt}$.²² The choice $j = 0$ represents the choice of treatment other than a statin, with mean utility normalised 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 versus alternative treatments, and ϵ_{idjt} is the standard type I extreme value error term (with scale normalised to one) that is i.i.d. across drugs. As the nesting parameter $\lambda \in [0, 1]$ approaches 1, there is less substitution outside the nest of statins.²³ 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 versus 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 favour 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:

$$\delta_{djt} = \theta_{dj}^m 1_{\{m_{dj} > 0\}} - \theta^p p_{djt}^{oop} + X_{djt} \theta_j^x + \zeta_{djt}. \quad (1)$$

Here, $1_{\{m_{dj} > 0\}}$ is an indicator for whether cardiologist d received a meal from j ’s manufacturer and θ_{dj}^m 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, $\theta^p p_{djt}^{oop}$ is the average OOP 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. $X_{djt} \theta_j^x$ 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, ζ_{djt} is a cardiologist-drug-year-level unobservable term, which we allow to have two components:

$$\zeta_{djt} = \tilde{\zeta}_{djt} + \varepsilon_{djt}^{de}. \quad (2)$$

22. 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.

23. 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.”

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

The decision error parameter approach has some appealing features. It can capture many theoretical frictions in a reduced form way (Baicker *et al.*, 2015; Mullainathan *et al.*, 2012). It is empirically flexible in that one can estimate decision utility following typical revealed preference-based procedures and then consider how different types of decision errors affect welfare. In prior studies with decision errors, data on *unbiased* decision-makers are leveraged to estimate true equilibrium welfare for the whole sample (Allcott and Taubinsky, 2015; Bronnenberg *et al.*, 2015; Handel and Kolstad, 2015). In Section 5, we discuss how outside data might be used to calibrate a decision error in (the many) empirical contexts such as ours where no unbiased decision-makers are identified. At this stage, we leave the decision error specification fully flexible in terms of the mean decision error, heterogeneity in errors 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 maximises decision utility, so that expected quantities demanded are given by:

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

Given this model, we represent expected consumer surplus as:

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

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 (versus 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- ($\varepsilon_{djt}^{de} < 0$) or over-prescribing ($\varepsilon_{djt}^{de} > 0$), as well as the countervailing (or reinforcing) effect of meals.^{24,25} See Appendix C.2 for further discussion.

24. A related (and not mutually exclusive) interpretation would be that physicians maximise a sum of physician (chooser) and patient (consumer) utility, with ε_{dj}^{de} governing the difference between the physician’s maximum and the patient’s.

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

3.2. Targeting meal payments to physicians

The parameter θ_{dj}^m 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:

$$1\{m_{dj} > 0\} = X_{dj}\gamma_j^x + Z_{dj}\gamma_j^z + \mu_{dj}. \quad (5)$$

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 AMC 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 conceptualise this as a simultaneous game of incomplete information where the competing firm's expected strategy is captured in $X_{dj}\gamma_j^x + Z_{dj}\gamma_j^z$, 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 do not 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 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 PBMs and insurers).²⁶

Let the supplier's profit be: $\pi(p_{jrt}^{pos}) = \sum_{d \in r} q_{djt}(p_{jrt}^{pos}(1 - \tau_{jt}) - mc_{jt})$, where τ_{jt} is the manufacturer rebate as a fraction of the POS price and mc_{jt} captures the cost of manufacturing

26. As discussed by [Starc and Swanson \(2021\)](#), 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.

and distributing the marginal unit of drug j in dollars.²⁷ p_{jrt}^{pos} is the POS price insurers pay for the drug, which we model as constant across cardiologists within region r . We link the negotiated POS price and OOP price paid by enrollees via $p_{djt}^{oop} = cs_{djt} p_{jrt}^{pos}$, 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.²⁸ We take the region r over which POS 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.²⁹

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.*, 2019). In the model, each price maximises the Nash Product of the gains from trade for each supplier and buyer pair, taking other prices as given. The first-order conditions of this model (see Appendix C for details) generate pricing equations that can be represented by:

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

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 ($b_{jrt} = 1$) versus 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 OOP price p_{djt}^{oop} given insurance coverage, but the insurer and supplier negotiate over POS price p_{jrt}^{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

27. 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.

28. 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 favourably.

29. 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.

the relative weight insurers place on enrollee surplus and plan costs:³⁰

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

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.³¹

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.³² 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 linearise 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: $\delta_{djt} = \ln(s_{djt}/s_{d0t}) - \lambda \ln(s_{djl_{gt}})$. Combining this with equation (1) yields the linear specification:

$$\ln(s_{djt}/s_{d0t}) = \lambda \ln(s_{djl_{gt}}) - \theta^p p_{djt}^{oop} + \theta_{dj}^m 1_{\{m_{dj} > 0\}} + X_{djt} \theta_j^x + \zeta_{djt}. \quad (7)$$

30. 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.

31. 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 favourable 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 (Olssen and Demirer, 2019; Dafny *et al.*, 2024) on that process.

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

where s_{djt} is j 's overall market share, s_{d0t} is the market share of the outside good (nonstatin treatments), and $s_{dj|g_t}$ is j 's market share within nest g , the set of statin treatments. During 2011–12, nonstatin 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 atorvastatin at the end of 2011 to identify the coefficients on price and within-nest share. We estimate:

$$\ln(s_{djt}/s_{d0t}) = \lambda \ln(s_{dj|g_t}) - \theta^p p_{djt}^{oop} + \psi_{dj} + \theta_t + \theta_{Lip12} + \xi_{djt} \quad (8)$$

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 versus the composite outside good of nonstatin 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_{djt}^{oop} by constructing IV 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, versus the generic statins from the earlier generation (described in detail in Appendix A.2).³³ 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 favourable 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 use 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 OOP 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: $p_{djt}^{oop, IV}$. We also create an analogous instrument based on an average dummy for formulary inclusion: $\bar{1}_{\{j \in form_{djt}^{IV}\}}$. We denote these two instruments: $Z^p = [p_{djt}^{oop, IV}, \bar{1}_{\{j \in form_{djt}^{IV}\}}]$. 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 mechanically affects within group shares, by adding polynomials in the cardinality of the sets

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

of statins and strong statins prescribed $Z^J = [\ln(|\mathcal{J}_{dt}|), |\mathcal{J}_{dt}|, |\mathcal{J}_{dt}|^2, \ln(|\mathcal{J}_{dt}^{ss}|), |\mathcal{J}_{dt}^{ss}|, |\mathcal{J}_{dt}^{ss}|^2]$ 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 dependent variable only through their effects on $(\ln(s_{dj|g,t}), p_{dj}^{oop})$.

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 the drug-doctor meal indicator and a rich set of controls for physician and market characteristics:

$$\hat{\psi}_{dj} = \theta_{dj}^m 1_{\{m_{dj}>0\}} + \theta_j + \bar{X}_{dj} \bar{\theta}_j^x + \bar{\xi}_{dj}. \quad (9)$$

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.³⁴ In our preferred specification, we construct $1_{\{m_{dj}>0\}}$ as a dummy for physician d receiving any payment from Pfizer over 2010–12 (in the case of $j = \text{Lipitor/atorvastatin}$), or as a dummy for physician d receiving any payment from AstraZeneca over 2011–12 (in the case of $j = \text{Crestor}$).³⁵ 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, 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 the average level of the unobserved preference of doctor d for drug j , $\bar{\xi}_{dj}$, and the heterogeneous component of the same doctor's expected treatment response for that drug, θ_{dj}^m . Accordingly, we develop an econometric specification that estimates the full distribution of treatment effects θ_{dj}^m across physicians. We estimate the distribution of 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

34. 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.

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

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 IV strategy is designed to recover estimates of MTEs as our continuous instrument varies ([Heckman and Vytlačil, 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 probability of receiving the meal treatment (as a function of the instrument, conditional on observables) becomes infinitesimally small; *i.e.* as we compare doctors in regions with closer CoI policies, conditional on observables. To implement this strategy, the `mtefe` package follows [Heckman and Vytlačil \(2007\)](#) in specifying the conditional expectation of the error term $\tilde{\epsilon}_{dj}$ as a nonparametric function of the treatment propensity given controls $\tilde{X}_{dj}$ 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. 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.

To summarise: We face a scenario where manufacturers’ decisions to interact with physicians may depend on physicians’ heterogeneous preferences over drugs and responses to interactions, both of which manufacturers can predict to some degree. We collect data on a large number of observable physician and regional characteristics that help predict interactions and prescribing. Still, heterogeneity that is unobservable to us (but observable to manufacturers and/or physicians) may remain. To address this, we construct an instrumental variable based on how AMC policies affect the likelihood of manufacturer-physician interactions across regions. We also show that, based on observables, the instrument appears to be effectively randomly assigned after conditioning on a few regional characteristics. The MTE framework explicitly allows for physicians’ responsiveness to treatment (identified by the variation generated by the instrument) to vary with their unobserved likelihood of being targeted for interactions, thus providing estimates of the distribution of responsiveness across physicians. Combined with physicians’ observable characteristics and actual interactions, this in turn allows us to estimate each physician’s expected responsiveness to interactions. Finally, we obtain an estimate of their unobserved preference for each drug as the residual from the prescribing model.

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 $2 \times \sqrt{N_d}$ 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 IV 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 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% 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 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

TABLE 3
Demand estimates step 1—price and nest coefficients and $\{\psi\}$

	(1)	(2)	(3)
θ^P	0.00173 (0.00004)	−0.00027 (0.00003)	−0.00816 (0.00018)
λ	0.943 (0.0005)	0.961 (0.0009)	0.339 (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}/ \theta^P $)strongstatis			−306.7
mean($\psi_{dj}/ \theta^P $)othergenerics			−327.1
s.d.($\psi_{dj}/ \theta^P $)			80.4

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

(a)

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




FIGURE 3

Explaining variation in meal payments (a) First stage regression and (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 standardised 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.

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 versus the unobserved resistance to treatment.³⁶ 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 10th 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

36. 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

MTEs of meal payments (a) treatment effect distribution and (b) treatment effect point estimates

Notes: Panel (a) plots the MTE curve ($E[\theta^m | U = u]$) with 95% 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.

decline), while at the 90th percentile of unobserved 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 MTEs.³⁷ 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 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 direction. This logic is outlined in greater detail in Appendix C.3, and Appendix D.2 relates our econometric specification to a generalised 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 generalised Roy framework (see Heckman and Vytlačil, 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 (Shapiro, 2018a; Agha and Zeltzer, 2019; Carey *et al.*, 2020). For example, Shapiro (2018a) finds

37. 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 IV methods are not guaranteed to weight MTEs 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.

that a detailing visit increases prescribing of antipsychotics by 14% in the subsequent year,³⁸ whereas the coefficients in our nested logit demand model imply that a meal-based relationship increases promoted statin prescribing by about 32% for the average physician, but by 54% for the average physician actually targeted by firms and only 8% 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.³⁹

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 realised treatment) to derive the expected response to treatment $E[\theta_{dj}^m]$ for any observation in the data.⁴⁰ Figure 5, Panel (a) presents a histogram of expected treatment effects, normalised by the standard deviation of the physician-molecule preference variation, $E[\theta_{dj}^m]/SD(\psi_{dj}^m)$, separately for physicians who in fact did or did not receive a meal payment. 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% of these not-paid cardiologists the effect is not statistically different from zero. The difference between the centres 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.

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.⁴¹ 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

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

39. Chintagunta and Manchanda (2004), Shapiro (2018a), and 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.

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

41. Liu *et al.* (2020) estimate that Pfizer (AstraZeneca) visited each detailed physician 9.79 (6.90) times per year in 2002–04 to discuss Lipitor (Crestor), at an estimated cost of \$150 (\$187) per visit in 2003 dollars. This implies a “cost of relationship” of about \$1,780.69 (\$1,563.65) per physician-year in 2011 dollars, before accounting for the \$50–\$80 direct cost of payments.




FIGURE 5

Heterogeneity in expected treatment effects across doctors, by actual treatment status (a) Prescribing Change with Payment (s.d.). (b) Prescribing without payments (share) and (c) Manufacturer return to payment (\$)

Notes: Each figure displays two histograms: one for the sample of physicians who did receive payments in the data and one for those who did not. Panels correspond to different variables summarised in this manner. 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 $E[s_{dj}^{m=0}]$; 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.

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

To estimate the model for a given vector (ε^{de} , τ_{jt} , mc_{jt}), we parameterise 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.⁴² 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 (Grennan, 2013; Gowrisankaran *et al.*, 2015). Our solution is to use estimates of rebates and marginal costs from recent research on pharmaceutical markets, and we test the sensitivity of our results to alternative assumptions. For example, in our baseline analysis, we assume that rebates for branded drugs were 26.3%, consistent with the average rebates for cardiovascular drugs reported to CMS in 2014, and we increase rebates to 48.3% 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% of the average POS price of generic statins: $mc = 0.17 * \overline{p_{gen}^{pos}}$. The value of 17% 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 summarises our supply side parameter estimates. The most striking feature is the high bargaining parameter estimates for the branded drugs relative to generics. Because the generic sales are aggregated over firms, the bargaining parameters also capture within-molecule competitiveness. This can also be seen in the slightly larger bargaining parameter for generic atorvastatin, where only two manufacturers compete during the first six months of 2012, after which eleven more manufacturers enter. The larger bargaining parameters for Lipitor and Crestor in 2012 reflect the fact that POS prices remain high in many regions for much of 2012 as insurers are slow to adjust formularies, despite the improved outside option with generic atorvastatin entry.⁴³

42. 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.

43. 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.

TABLE 4
Supply parameter estimates

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

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.⁴⁴

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 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 POS and OOP 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 OOP 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.

44. 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.

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 (0.001)	0.149 (0.002)	0.150 (0.002)	0.163 (0.002)	0.174 (0.001)
$Q_{Lipitor}$	0.034 (0.001)	0.029 (0.001)	0.029 (0.001)	0.040 (0.002)	0.047 (0.001)
$Q_{Crestor}$	0.022 (0.000)	0.016 (0.001)	0.016 (0.001)	0.023 (0.002)	0.031 (0.000)
$OOP_{statins}$	18.73 (0.21)	18.73 (0.21)	18.53 (0.22)	2.15 (0.1)	2.15 (0.1)
$POS_{statins}$	74.91 (0.74)	74.91 (0.74)	74.10 (0.73)	86.15 (1.43)	87.65 (1.32)

Notes: Authors' calculations of observed and counterfactual equilibrium statin quantities (share of cardiovascular utilisation) and prices (POS and OOP 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% points—a five percent reduction in total statin usage. For the focal branded statins, the decrease is 0.12% points, which is a 19% 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% points with meals (compare "Observed" to "No Ban, $p = mc$ ") and 1.3% 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 utilisation to undershoot the efficient allocation; in the Crestor case, meals cause utilisation to fall almost exactly at the allocation with $p = mc$ and a meal ban.

The effect of a ban on utilisation 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 POS 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 OOP prices from POS prices, POS prices would increase by 15–17%.

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 underutilisation. Motivated by the American College of Cardiology's position that statins are underused 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 OOP prices. $CS_{dt}(\mathcal{J}_t) - \sum_j q_{djt}(p_{jrt}^{pos}(1 - \tau_{jt}) - p_{djt}^{oop})$ (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 subsidised 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: $PS_{jrt} = \sum_{d \in r} q_{djt}(p_{jrt}^{pos}(1 - \tau_{jt}) - mc_{jt}) - C_{jr}^{m_{dj}=1}$, 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.⁴⁵ 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.

The results are summarised 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

45. As discussed previously, we assume that marginal manufacturing costs are 17% 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 the 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.

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 $\varepsilon^{de} = 0$).⁴⁶ 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 the 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 under-utilisation 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 we estimate here is fairly large, with a meal payment ban resulting in an approximately 15% 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

46. 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, and (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. The indicated value of ε^{LYG} is the decision error implied by our demand model in Section 3.1, if one assumes that the average true benefit of statins is equal to the dollar value of the clinical health benefits of statins as calculated in Appendix F.

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-versus 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 internalise 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 welfare 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.

Our results also shed some light on the value of the AMC CoI policies that generate our identifying variation. But while our findings can speak to the welfare impact of those policies specifically within the context of AMC physicians' prescribing behaviours for drugs similar to the statins we focus on, we caution from making any general conclusions about the overall net value of those policies given the numerous other factors at play. Perhaps foremost, welfare effects may vary across drugs. More broadly, CoI policies affect interactions beyond detailing. For example, a unique feature of AMCs is their central role in clinical trials and the dissemination of information about new therapies (Agha and Molitor, 2018). Many CoI policies are focused on physician-firm relationships in that context, and there is evidence that firm sponsorship of a clinical trial can have a significant influence on its results (Oostrom, 2023), which suggests a unique channel for CoI policies at AMCs to be valuable.

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 underutilisation (absent meals) surely varies across drugs. Many studies of statins point to potential underutilisation 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 “behavioural 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.⁴⁷ This perspective is also consistent with our result that meals tend to bring otherwise

47. Statins are far from the only example of underutilisation 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

low prescribers closer to the prescribing behaviour 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 nonadherent patients, and the dollar value of a life-year gain (see Appendix F for details), this implies that the decision-maker optimising on behalf of an average indicated patient should compare the monthly OOP price of statins (around \$33, in the case of branded statins Crestor and Lipitor) to a “flow” willingness-to-pay value of \$516.⁴⁸ This measure suggests that the $\varepsilon^{de}/\theta^p$ consistent with statin demand in our data 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—U.K. 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 utilisation by 5%, and the American College of Cardiology indicates that utilisation of statins should *increase* by 24% 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% of (randomly selected) Medicare patients in our sample would experience the clinical benefits of statins from the medical literature, even if the remaining 39% 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.

on subsequent patient health outcomes, consistent with health benefits being dramatically underweighted in decision utility.

48. 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 OOP costs that is necessary to cover five years of prescription statins, in present discounted value.

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 trade-offs, 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 IV, showing that local AMC CoI 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 MTEs 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 behaviour 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.

Acknowledgments. 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 (U.S.), 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 analysing and preparing the results reported herein. This paper subsumes an earlier working paper titled “Physician-Industry Interactions: Persuasion and Welfare”.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Data Availability Statement

The data and code underlying this research is available on Zenodo at: <https://dx.doi.org/10.5281/zenodo.11551801>.

REFERENCES

- 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 MOLITOR, D. (2018), “The Local Influence of Pioneer Investigators on Technology Adoption: Evidence from new Cancer Drugs”, *Review of Economics and Statistics*, **100**, 29–44.
- 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., *et al.* (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., *et al.* (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**, 1–15.
- ANDRESEN, M. E. (2018), “Exploring Marginal Treatment Effects: Flexible Estimation Using Stata”, *The Stata Journal*, **18**, 118–158.
- ARCIDIACONO, P., ELLICKSON, P. B., LANDRY, P., *et al.* (2013), “Pharmaceutical Followers”, *International Journal of Industrial Organization*, **31**, 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**, 1623–1667.
- BELLONI, A., CHERNOZHUKOV, V., FERNANDEZ-VAL, I., *et al.* (2017), “Program Evaluation and Causal Inference with High-Dimensional Data”, *Econometrica*, **85**, 233–298.
- BERRY, S. T. (1994), “Estimating Discrete-Choice Models of Product Differentiation”, *RAND Journal of Economics*, **25**, 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**, 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**, 985–1039.
- BRONNENBERG, B. J., DUBÉ, J. -P., GENTZKOW, M., *et al.* (2015), “Do Pharmacists buy Bayer? Informed Shoppers and the Brand Premium”, *Quarterly Journal of Economics*, **130**, 1669–1726.
- BROT-GOLDBERG, Z. C., CHANDRA, A., HANDEL, B. R., *et al.* (2017), “What Does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics”, *The Quarterly Journal of Economics*, **132**, 1261–1318.
- BROWN, Z. Y. (2019), “An Empirical Model of Price Transparency and Markups in Health Care” (Working Paper).
- CAMPBELL, J. J. (2008), *Understanding Pharma* (2nd edn) (Raleigh, NC: Pharmaceutical Institute).
- CARDELL, N. S. (1997), “Variance Components Structures for the Extreme-Value and Logistic Distributions with Application to Models of Heterogeneity”, *Econometric Theory*, **13**, 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**, 1299–1342.
- CHANDRA, A., FINKELSTEIN, A., SACARNY, A., *et al.* (2016), “Health Care Exceptionalism? Performance and Allocation in the us Health Care Sector”, *American Economic Review*, **106**, 2110–2144.

- 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**, 193–213.
- CHERNOZHUKOV, V., CHETVERIKOV, D., DEMIRER, M., *et al.* (2018), "Double/debiased Machine Learning for Treatment and Structural Parameters", *Econometrics Journal*, **21**, C1–C68.
- CHINTAGUNTA, P. and MANCHANDA, P. M. (2004), "Responsiveness of Physician Prescription Behavior to Salesforce Effort: An Individual Level Analysis", *Marketing Letters*, **15**, 129–145.
- CHOUDHRY, N. K., AVORN, J., GLYNN, R. J., *et al.* (2011), "Full Coverage for Preventive Medications after Myocardial Infarction", *New England Journal of Medicine*, **365**, 2088–2097.
- CLEMENS, J. and GOTTLIEB, J. D. (2014), "Do Physicians' Financial Incentives Affect Medical Treatment and Patient Health?", *American Economic Review*, **104**, 1320–1349.
- CMS (2012), "Centers for Medicare and Medicaid Services: Research, Statistics, Data and Systems".
- (2013), "Centers for Medicare and Medicaid Services: Provider Data".
- COLLARD-WEXLER, A., GOWRISANKARAN, G. and LEE, R. S. (2019), "Nash-In-Nash Bargaining: A Microfoundation for Applied Work", *Journal of Political Economy*, **127**, 163–195.
- 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**, 643–685.
- CURRIE, J. M. and MACLEOD, W. B. (2020), "Understanding Doctor Decision Making: The Case of Depression Treatment", *Econometrica*, **88**, 847–878.
- CUTLER, D. M., SKINNER, J. S., STERN, A. D., *et al.* (2019), "Physician Beliefs and Patient Preferences: A new Look at Regional Variation in Health Care Spending", *American Economic Journal: Economic Policy*, **11**, 192–221.
- DAFNY, L., HO, K. and KONG, E. (2024), "How Do Copayment Coupons Affect Branded Drug Prices and Quantities Purchased?", *American Economic Journal: Economic Policy*, **16**, 314–346.
- DATTA, A. and DAVE, D. (2016), "Effects of Physician-Directed Pharmaceutical Promotion on Prescription Behaviors: Longitudinal Evidence", *Health Economics*, **26**, 450–468.
- 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**, 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**, 1712–1752.
- DEJONG, C., AGUILAR, T., TSENG, C.-W., *et al.* (2016), "Pharmaceutical Industry-Sponsored Meals and Physician Prescribing Patterns for Medicare Beneficiaries", *JAMA Internal Medicine*, **176**, 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**, 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**, 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**, 122–153.
- FENG, J. (2022), "History-dependence in Drug Demand: Identification and Implications for Entry Incentives", *The Review of Economics and Statistics*, **106**, 455–469.
- FUGH-BERMAN, A. and AHARI, S. (2007), "Following the Script: How Drug Reps Make Friends and Influence Doctors", *PLoS Med*, **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**, 61–69.
- GOWRISANKARAN, G., NEVO, A. and TOWN, R. (2015), "Mergers When Prices are Negotiated: Evidence from the Hospital Industry", *American Economic Review*, **175**, 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**, 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**, 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**, 459–665.

- HANDEL, B. R. and KOLSTAD, J. (2015), "Health Insurance for 'humans': Information Frictions, Plan Choice, and Consumer Welfare", *American Economic Review*, **105**, 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**, 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 15 January 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**, 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**, 65–72.
- HECKMAN, J. J., URZUA, S. and VYTLACIL, E. (2006), "Understanding Instrumental Variables in Models with Essential Heterogeneity", *Review of Economics and Statistics*, **88**, 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**, 379–417.
- HO, K. and PAKES, A. (2014), "Hospital Choices, Hospital Prices, and Financial Incentives to Physicians", *American Economic Review*, **104**, 3841–3884.
- IIZUKA, T. (2012), "Physician Agency and Adoption of Generic Pharmaceuticals", *American Economic Review*, **102**, 2826–2858.
- INDERST, R. and OTTAVIANI, M. (2012), "Competition Through Commissions and Kickbacks", *American Economic Review*, **102**, 780–809.
- JACOBSON, M., O'MALLEY, A. J., EARLE, C. C., *et al.* (2006), "Does Reimbursement Influence Chemotherapy Treatment for Cancer Patients", *Health Affairs*, **25**, 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." (Working Paper No. 26846, NBER).
- KESSELHEIM, A. S., AVORN, J. and SARPATWARI, A. (2016), "The High Cost of Prescription Drugs in the United States", *JAMA*, **316**, 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., *et al.* (2008), "Generalizations on the Effectiveness of Pharmaceutical Promotional Expenditures", *International Journal of Research in Marketing*, **25**, 234–246.
- LARKIN, I., ANG, D. and STEINHART, J. (2017), "Association between Academic Medical Center Pharmaceutical Detailing Policies and Physician Prescribing", *JAMA*, **317**, 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**, 599–611.
- LIU, Q., GUPTA, S., VENKATARAMAN, S., *et al.* (2020), "An Empirical Model of Drug Detailing: Dynamic Competition and Policy Implications", *Management Science*, **62**, 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**, 785–822.
- MANINGAT, P., GORDON, B. and BRESLOW, J. (2013), "How Do We Improve Patient Compliance and Adherence to Long-Term Statin Therapy?", *Current Atherosclerosis Reports*, **15**, 291–291.
- MCFADDEN, D. (1978), "Modeling the Choice of Residential Location", *Transportation Research Record*, **673**, 72–77.
- 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**, 69–85.
- NARAYANAN, S. and MANCHANDA, P. (2009), "Heterogeneous Learning and the Targeting of Marketing Communication for New Products", *Marketing Science*, **28**, 424–441.
- OLSSSEN, A. L. and DEMIRER, M. (2019), "Drug Rebates and Formulary Design: Evidence from Statins in Medicare Part D" (Mimeo).
- OOSTROM, T. (2023), "Funding of Clinical Trials and Reported Drug Efficacy" (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**, 406–433.
- SCOTT MORTON, F. and KYLE, M. (2012), "Markets for Pharmaceutical Products", in McGuire, T. (ed) *Handbook of Health Economics* (Vol. 2, Chapter 12) (Amsterdam: Elsevier B.V.).

- SHAPIRO, B. (2018a), “Informational Shocks, Off-Label Prescribing and the Effects of Physician Detailing”, *Management Science*, **64**, 5925–5945.
- (2018b), “Positive Spillovers and Free Riding in Advertising of Prescription Pharmaceuticals: The Case of Antidepressants”, *Journal of Political Economy*, **126**, 381–437.
- SINKINSON, M. and STARC, A. (2019), “Ask Your Doctor? Direct-To-consumer Advertising of Pharmaceuticals”, *The Review of Economic Studies*, **86**, 836–881.
- SOOD, N., SHIH, T., NUYS, K. V., *et al.* (2017), “The Flow of Money Through the Pharmaceutical Distribution System” USC Schaeffer Center.
- SPURLING, G. K., MANSFIELD, P. R., MONTGOMERY, B. D., *et al.* (2010), “Information from Pharmaceutical Companies and the Quality, Quantity, and Cost of Physicians’ Prescribing: A Systematic Review”, *PLoS Med*, **7**, e1000352.
- STARC, A. and SWANSON, A. (2021), “Preferred Pharmacy Networks and Drug Costs”, *American Economic Journal: Economic Policy*, **13**, 406–446.
- TARN, D. M., PLETCHER, M. J., TOSQUI, R., *et al.* (2021), “Primary Nonadherence to Statin Medications: Survey of Patient Perspectives”, *Preventive Medicine Reports*, **22**, 133–139.
- TUNCEL, T. (2024), “Should Off-Label Drug Prescriptions be Prevented? Empirical Evidence from France”, *Review of Economic Studies*, **92**, 2027–2070.
- VAN DULMEN, S., SLUIJS, E., VAN DIJK, L., *et al.* (2007), “Patient Adherence to Medical Treatment: A Review of Reviews”, *BMC Health Services Reviews*, **4**, 1–16.
- 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., *et al.* (2016), “Association of Industry Payments to Physicians with the Prescribing of Brand-Name Statins in Massachusetts”, *JAMA Internal Medicine*, **176**, 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?”