American Economic Journal: Applied Economics 2020, 12(3): 1–39
https://doi.org/10.1257/app.20170647

The Impact of Information Technology on the Diffusion of New Pharmaceuticals[†]

By Kenneth J. Arrow, L. Kamran Bilir, and Alan Sorensen*

Do information differences across US physicians contribute to treatment disparities? This paper uses a unique new dataset to evaluate how changes in physician access to a decision-relevant drug database affect prescribing decisions. Our results indicate doctors using the reference have a significantly greater propensity to prescribe generic drugs, are faster to begin prescribing new generics, and prescribe a more diverse set of products. These results are consistent with database users responding primarily to the increased accessibility of non-clinical information such as pricing and insurance formulary data, and suggest improvements to physician information access have important implications for aggregate healthcare costs. (*JEL* D83, I11, I18, L65)

National health expenditures exceed $3 trillion annually in the United States, account for nearly 20 percent of US GDP, and are to a considerable extent publicly funded.[1] Yet, research by the Dartmouth Atlas Project and Cooper et al. (2015) finds substantial, systematic disparities in both the extent of health spending and the quality of medical care across US regions, including threefold per capita expenditure gaps resulting from inefficient variations in care—differences consistent neither with patient preferences nor with underlying medical conditions.[2] These findings imply significant gains could be achieved by improving efficiency in low-performing regions, but this requires first identifying the specific mechanisms that *cause* treatment disparities. Among the many potential mechanisms that have been proposed—which include supply, demand, regulatory, and pricing differences—perhaps the most important and intriguing is that disparities result from a

*Arrow: Formerly Department of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94305 (email: arrow@stanford.edu); Bilir: Department of Economics, University of Wisconsin-Madison, 1180 Observatory Drive, Madison, WI 53706 (email: kbilir@ssc.wisc.edu); Sorensen: Department of Economics, University of Wisconsin-Madison, 1180 Observatory Drive, Madison, WI 53706 (email: sorensen@ssc.wisc.edu). Dave Donaldson was coeditor for this article. We thank the Ewing Marion Kauffman Foundation for its financial support of this research; Bilir also thanks the Princeton University IES for its hospitality and support. We thank Murray Aitken, Nick Bloom, Michael Dickstein, Jeff Gambino, Shane Greenstein, Ben Handel, Bob Litan, Greg Rosston, Randall Stafford, Jeff Tangney, Lee Vermuelen, the UW Health Physician Informatics Team, and the research team at IMS Health (IQVIA) for guidance and very helpful conversations.

[†]Go to https://doi.org/10.1257/app.20170647 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

[1]See https://www.cms.gov/nationalhealthexpenddata.

[2]The Dartmouth Atlas Project has documented healthcare disparities for Medicare patients over decades (Wennberg and Cooper 1996; Gawande 2009; Chernew, Hirth, and Cutler 2009). Cooper et al. (2015) find related disparities among the privately insured population. See also Wennberg and Wennberg (2003).

lack of uniformity in physicians' information about available therapies.[3] The possibility that information problems underlie observed treatment disparities has inspired calls for the expanded use of medical decision aids, but the difficulties inherent to measuring information differences have led to a paucity of systematic evidence on their actual importance.[4]

To shed light on the question of how physicians' information access impacts treatment choices, we assemble a new dataset in which treatment decisions and access to a decision-relevant database are directly observed at the doctor level for the universe of US prescribers over the decade 2000–2010.[5] Our ability to observe both treatment decisions and database usage by individual doctors over a period of time is a particularly important and unique aspect of the dataset we construct; furthermore, these data cover Medicare patients and the privately insured as well as those with Medicaid or no coverage, so that the results of our analysis are representative of the full range of US patient types.[6]

Using these data, we provide novel evidence that physicians using the drug reference database significantly increase their likelihood of prescribing a generic drug relative to brand name therapies, and thus increase their generic prescription share—one of the key efficiency metrics emphasized in the Dartmouth Atlas in documenting prescription drug variations.[7] Database users are also significantly faster to begin prescribing newly released generics, an effect absent for new branded drugs. These findings suggest that database users may be responsive to the increased salience of non-clinical information in the database—including whether a particular drug is currently covered by a patient's insurance plan and plan-specific pricing—as a generic drug and its branded equivalent share essentially identical clinical attributes. We find that treatment differences across doctors decline significantly more among database users than nonusers during the sample period, while the actual diversity of a user's own prescribing increases on average following adoption. Access to detailed information about competing treatments thus appears to raise efficiency and reduce disparities, but importantly, these effects do not appear to come at the expense of patient-sensitive decision making.

Our empirical strategy is data intensive, requiring the combination of two unusually large proprietary datasets; therefore, we focus our analysis on a single class of pharmaceutical products—cholesterol drugs.[8] While this drug class is already of immediate interest due to its exceptionally large US market, the rapid pace of innovation during the period relevant for doctors' adoption of the drug database imply this class is ideal for studying the effects of information technology usage on the diffusion of new prescription drugs. Indeed, during the relevant decade, twelve

[3] For example, Skinner (2011).

[4] See Phelps (2000), Wennberg and Wennberg (2003), and Arrow (1969).

[5] The provider of the database is a leading US point-of-care medical applications firm that chose to remain unnamed in this study. For a description of all major drug references, see Ventola et al. (2014).

[6] These categories are new relative to existing evidence on healthcare variations; see footnote 2.

[7] See Munson et al. (2013).

[8] Prescription data covering additional drug classes exist, but are not available for the current study due to the unusually large size of the customized data extracts involved.

nationwide product innovations occurred.[9] The differential response across physicians to these repeated drug introductions is crucial to our identification strategy because it allows us to measure the influence of information access on treatment decisions while accounting for physician characteristics that simultaneously affect both prescribing and database adoption.

This is particularly important for our analysis because access to the drug database is not randomly assigned—doctors choose whether and when to subscribe. Prescription patterns of subscribing doctors may therefore look different from those of nonsubscribers not due to any effects of the database itself, but instead due to differences in the types of doctors who choose to subscribe. With this challenge in mind, our analysis relies heavily on within-doctor variation over time: rather than estimating effects by comparing database users to nonusers, we focus on comparisons of a doctor's own prescriptions before versus after she begins using the database. In line with this strategy, much of our main analysis restricts attention to the sample of physicians that eventually adopts the reference database. To account for the possibility that dynamic prescribing determinants may be correlated with adoption timing within a location, such as changes in doctor-specific drug advertising, our main specifications also include doctor-specific time trends and zip-code-month fixed effects. We further provide time-varying estimates that indicate the prescribing changes we find either coincide with or immediately follow database adoption.

Our empirical approach nevertheless leaves open the possibility that unobserved changes not accounted for by doctor-specific time trends or zip-code-month fixed effects both affect prescribing and coincide with database adoption. To help address this concern, we group doctors based on the intensity with which the drug database is used to search about cholesterol drugs. In line with the idea that doctors' prescribing changes are associated with use of the database, these results reveal a larger association between database usage and prescribing changes among relatively intense database users, compared with lower-intensity users. Also consistent with this interpretation, the drug reference database we consider is a stand-alone technology, accessed by an individual rather than institutional subscription, and is not embedded into other health information technology systems. This is important, because if the database were part of a broader health information technology platform—for example, one that includes features that facilitate patient and insurer billing (see Agha 2014)—it would be difficult to know whether estimates corresponding to database use indicate effects of drug information access or, instead, effects of some other technology that is part of the same health IT system that the doctor accesses.[10]

[9] The Centers for Disease Control and Prevention estimates that approximately 71 million US adults suffer from chronic hypercholesterolemia and dyslipidemia, conditions in which abnormal levels of cholesterol or lipids are present in the bloodstream. These conditions are associated with heart disease, heart attack risk, and premature death; accordingly, sales of cholesterol therapies accounted for over $18 billion in 2011 (Ledford 2013). See also Mozaffarian et al. (2015). Drug introductions are described in Section IIA and listed in Table 1.

[10] Relatedly, it is important to note that while other drug references exist, including the Micromedex and UpToDate Lexicomp databases, these are imperfect substitutes for the database we consider. In particular, these other databases do not contain drug price or formulary-specific coverage information and are typically accessed by institutional subscription. Because our data do not include information on doctors' potential usage of these alternative platforms, it is also important to point out that if a doctor tends to begin using all databases around the

Consistent with the regional disparities documented in the Dartmouth Atlas, our data reveal substantial prescribing variation across the universe of individual US physicians, particularly with respect to generics and new drug adoption. Doctors differ widely in generic prescription shares (mean 56.4 percent, standard deviation 24.3 percent) in December 2010, and span the full range from no generics (fifth percentile) to only generics (ninety-fifth percentile).[11] We find that these generic shares are strongly and positively correlated within physician across patient insurance types (e.g., private versus Medicare), suggesting patient cost sharing is unlikely to explain the observed heterogeneity in prescribing patterns. Moreover, with respect to drug adoption, some US physicians begin prescribing a newly-approved cholesterol drug immediately, while others delay for a year or more before prescribing it, a pattern strongly evident even among new generic drugs. Six months after the introduction of generic lovastatin, for example, the molecular equivalent of Mevacor, the generic version accounted for only 83 percent of the drug's prescriptions; by contrast, the generic share reached essentially 100 percent by December 2010. This delayed substitution is evident for each of the generic entrants we observe, contributing to wide differences across doctors in the overall prescription share of generic drugs and to large corresponding gaps in cost.[12]

Our empirical analysis indicates that some of this observed prescription heterogeneity is explained by differences across doctors in information access. Our most conservative estimates indicate that, after obtaining database access, a physician user increases the likelihood of prescribing a new generic drug within its first market year by 1.3 percent; among high-intensity database users, this rises to a 2.4 percent increase. Regarding diversity, database users increase the number of unique drugs prescribed each month by a modest but highly significant 0.035 drugs, reducing the prescription Herfindahl-Hirschman index by 0.003 points. In line with both results, users increase the monthly likelihood of prescribing new and old generic drugs by 1.6 and 2.4 percent, respectively, after database adoption, while reducing the likelihood of prescribing a new branded drug by 0.5 percent. Our back-of-the-envelope calculations suggest the resulting increase in users' generic prescription shares contributes to substantial aggregate cost savings, which amount to approximately $1 billion annually for prescription drugs alone.[13]

This paper is related to an extensive literature documenting wide healthcare disparities across US regions, including the Dartmouth Atlas (e.g., Wennberg and Cooper 1996) and its analysis of prescription drug use among Medicare patients (Munson et al. 2013), and Cooper et al. (2015) for their analysis of the privately insured. We contribute to this work by first documenting prescription disparities for US prescribers and patients with all insurance types within a major therapeutic area. Second, relying on the unusual level of detail and coverage in the dataset we have assembled, we identify a highly significant link between observed disparities and a specific mechanism—physician information differences—that we find is

same time, our estimates would reflect the influence of information from all platforms rather than just the one we observe directly.

[11] See Table 3.

[12] See Section IID and Section VII.

[13] See Section VII.

partially responsible for these disparities. While our data are broader with respect to physician and patient coverage, our empirical strategy is demanding (our dataset includes over 200 million observations); we therefore focus on a single clinical area while Dartmouth Atlas and Cooper et al. (2015) cover a comprehensive set of treatments. Aggregating our physician-level data to Dartmouth Atlas regions, we nevertheless find that locations with high generic prescription shares in our dataset also have high generic shares for Medicare overall (correlation 44.4 percent), as well as lower per capita medical spending for prescription drugs (correlation 23.9 percent) and non-drug healthcare (correlation 12.4 percent)—patterns that underscore the highly systematic nature of US disparities in care, and suggest the potential value of extending our physician-level analysis to other clinical settings.[14]

Our paper is also closely related to work aimed at evaluating the impact of information technology on economic decisions and outcomes.[15] Agents' electronic information access can affect productivity (Solow 1987) and has been specifically shown to improve performance in emergency healthcare delivery (Athey and Stern 2002). However, in routine medical contexts the evidence is less clear. Dranove et al. (2014) find that the adoption of electronic medical records (EMR) raised hospital costs on average, with an important exception—adopting locations with an abundance of industrial IT did in fact experience cost declines. Our results complement this latter finding in that the medical decision support tool we study is standardized, likely to a greater extent than EMR, yet we observe that both the intensity and impact of its use differ substantially across physicians in the data. In particular, we find that the efficiency impact of database use is systematically larger among adopters using the database intensively to search for information about the cholesterol drugs we study.

In finding that use of an information database tilts prescribing away from branded drugs and toward generics, our results contribute to important work highlighting the influence of information on tastes for generic products. Bronnenberg et al. (2015) find that relatively informed buyers are more likely to choose a generic version, for example, when purchasing an over-the-counter drug, suggesting consumer misinformation contributes to the brand premium for health products. Our results add nuance to this finding, suggesting that even among highly trained and educated US physicians, access to current product information including pricing *increases* the propensity to prescribe a generic version and decreases that for branded drugs. We find that the impact of database access is systematically larger for physicians located far from the information frontier, and that dynamics in the product space (drug entry) may be important in explaining our results, as database users are also faster to begin prescribing a newly-introduced generic version. In finding that physicians' information access affects decisions made on behalf of patients, and that prescribing is highly correlated within a physician across patients regardless of insurance

[14] Our data do not include individual patient characteristics; this latter aspect precludes a direct extension of Munson et al. (2013) to non-Medicare patients, as well as a quantitative welfare analysis. It further precludes estimating a model featuring prescription dynamics within each patient-physician pair, as in Crawford and Shum (2005) or Dickstein (2018).

[15] See, for example, Attewell (1992); Bresnahan and Greenstein (1996); Black and Lynch (2001); Bresnahan, Brynjolfsson, and Hitt (2002); Brynjolfsson and Hitt (2003); Hubbard (2003); Forman, Goldfarb, and Greenstein (2005); Bloom et al. (2009); Bloom, Sadun, and Van Reenen (2012); Agha (2014).

coverage, our results are further aligned with Brot-Goldberg et al. (2015) and Cutler et al. (2015), which find evidence that physician preferences are key in explaining treatment decisions.

Our results add to the literature examining the determinants of new medical technology diffusion. Classic work by Coleman, Katz, and Menzel (1957, 1966) finds that new pharmaceutical products diffuse unevenly across medical practitioners: physicians who interact more frequently with other physicians are more likely to adopt early. Relatedly, Skinner and Staiger (2007) provide evidence that certain US states have a systematic tendency to adopt early across technology types as varied as beta blockers and hybrid corn (Griliches 1957). We find that physician access to a digital database also speeds new drug adoption, but only for generics; further, to account for local differences in the tendency to adopt both drugs and the database early, which could reflect general differences in unobserved factors such as drug advertising, we emphasize specifications that include both physician and zip-code-by-month fixed effects.[16] In focusing on individual-level drug adoption, our work is also closely related to Crawford and Shum (2005) and Dickstein (2018), who estimate models of physician learning, and Agha and Molitor (2018), who study the diffusion of cancer drugs.[17]

More broadly, our analysis complements research on general theories of technology diffusion featuring agents with imperfect information. Such theories can be shown to explain large existing differences in productivity across locations (Solow 1956, Arrow 1969, Parente and Prescott 1994, Comin and Hobijn 2004) as identified in Klenow and Rodríguez-Clare (1997) and Caselli and Coleman (2006), for example.[18] We introduce a unique dataset in which a sequence of technology adoption decisions is clearly observed at the individual level for the universe of US prescribers, allowing our study to speak both to microlevel mechanisms driving diffusion and to the aggregate consequences of these mechanisms.

The rest of the paper is organized as follows. Section I describes the data used in our analysis. Section II describes a simple model of prescription choice and our estimation framework. Section III presents the empirical results; Sections IV and V discuss interpretation; and Section VI concludes.

I. Data and Descriptive Evidence

Evaluating the influence of information access on new pharmaceutical drug diffusion requires detailed measures of drug innovations and individual prescribers' treatment decisions, information usage, and characteristics. We introduce each of these measures below and go on to describe physicians' prescribing of new and existing pharmaceutical drugs.

[16] We also consider the influence of local differences in mandatory substitution regulations that could be particularly important for explaining generic diffusion in the data, and find that the effects of database access are evident among physicians practicing both within and outside states with a mandatory substitution law.

[17] See also Escarce (1996), who studies physicians' decisions to adopt a surgical technology.

[18] The idea that underlying heterogeneity across agents could influence technology diffusion also relates our work to neoclassical models of technology adoption—e.g., David (1966) and Manuelli and Seshadri (2014).

A. *US Innovations in Chronic Hypercholesterolemia and Dyslipidemia Therapy*

Our analysis considers a class of prescription drugs—cholesterol medications—that are interesting both because of their substantial US market and, importantly, because of significant drug innovations during the time period relevant to physicians' adoption of the drug reference database. Specifically, at the start of our sample period in January 2000, HMG-CoA reductase inhibitors (statins) were understood to be the most effective pharmaceutical therapies for hypercholesterolemia, and there were five such products available: Lescol, Lipitor, Mevacor, Pravachol, and Zocor.[19] The most common non-statin used to treat high cholesterol was Niaspan, which is also included in our sample. Thereafter, twelve new cholesterol or lipid control therapies were introduced, including new formulations, combinations, and versions.[20] These include (i) three new molecular entities: Crestor, Lovaza, and Zetia; (ii) three generic versions: lovastatin (Mevacor), pravastatin (Pravachol), and simvastatin (Zocor); (iii) two new formulations: Altoprev (extended-release Mevacor) and Lescol XL (extended-release Lescol); and (iv) four new drug combinations: Advicor (extended-release niacin and Mevacor), Pravigard PAC (aspirin and Pravachol), Vytorin (Zetia and Zocor), Simcor (extended-release niacin and Zocor). Each new therapy received nationwide approval by the US Food and Drug Administration (FDA) on a known, drug-specific date (Table 1). All products are described in online Appendix A.1.

While these 18 products are therapeutic substitutes in that they aim at a similar clinical endpoint—cholesterol or trigliceride reduction—they are only imperfect substitutes: each product features distinctive characteristics relevant for the prescribing decision. First, many but not all cholesterol therapies are pure statins, which act to reduce cholesterol synthesis in the liver by inhibiting a specific coenzyme; these include Lescol (fluvastatin), Lipitor (atorvastatin), Mevacor (lovastatin), Pravachol (pravastatin), Zocor (simvastatin), Crestor (rosuvastatin), Altoprev (extended-release lovastatin), and Lescol XL (extended-release fluvastatin). Other products rely on different mechanisms of action. Zetia (ezetimibe), for example, is distinct in that it achieves cholesterol reduction by reducing intestinal absorption of cholesterol. A second distinction involves therapeutic intensity. High doses of Lipitor and Crestor are more effective at lowering low-density lipoprotein (LDL) cholesterol than alternatives (Law, Wald, and Rudnicka 2003). Side effects are also relevant; evidence suggests, for example, that high doses of Lipitor and Crestor may increase the incidence of adverse reactions, while combination therapies such as Vytorin may in certain cases be more appropriate care for patients with severe cholesterol abnormalities (Kastelein et al. 2008).

More subtly, clinical evidence suggests the benefits and risks associated with statins are heterogeneous across patients. Randomized controlled trials (RCTs) indicate, for example, that the benefits of statin use are higher for patients with

[19] Cannon et al. (2004).

[20] To ensure adequate coverage in the data, we consider all cholesterol therapies introduced by December 2008, but not those introduced after this date. For the same reason, our analysis excludes Baycol, a drug that was available in January 2000 but withdrawn from the market in August 2001.

Table 1—Descriptive Statistics, US Cholesterol Drug Introductions, January 2000–December 2008

Drug name	Release date	FDA approval category	Months to first prescription, conditional on prescription: Mean (1)	Months to first prescription, conditional on prescription: SD (2)	Adoption share December 2010 (3)
Lescol XL	October 2000	Dosage form	28.89	23.56	0.620
Advicor	December 2001	Combination	64.77	15.38	0.295
lovastatin	December 2001	Generic version	19.87	22.66	0.923
Altoprev	June 2002	Dosage form	42.62	24.07	0.151
Zetia	October 2002	Molecular entity	15.13	17.34	0.928
Pravigard PAC	June 2003	Combination	7.30	5.94	0.037
Crestor	August 2003	Molecular entity	22.67	21.83	0.923
Vytorin	July 2004	Combination	13.10	13.48	0.891
Lovaza	November 2004	Molecular entity	34.98	17.08	0.659
pravastatin	April 2006	Generic version	7.41	12.30	0.909
simvastatin	June 2006	Generic version	3.05	7.33	0.982
Simcor	February 2008	Dosage form	12.33	9.18	0.230

Notes: This table summarizes the variation across individual US physicians in the initial prescription of twelve new pharmaceutical products, each aimed at controlling blood cholesterol or lipid levels. Each product was approved for sale in the United States on the date indicated. New drug approvals are categorized by the FDA based on whether the product is a new molecular entity, a new drug combination, a new dosage form, or a new generic equivalent. The distribution of initial prescription dates across the set of US physicians that prescribe the drug at least once by December 2010 is described by the mean (column 1) and standard deviation (column 2) in months from initial FDA approval to the first prescription filled at a US pharmacy. The share of physicians that prescribe the product at least once by December 2010 (column 3) ranges from 3.7 percent (Pravigard PAC) to 98.2 percent (simvastatin).

Source: Prescription data are from IMS Health (IQVIA).

diabetes, negligible among those with prior heart failure, and vary with age; risks and side effects also vary with statin intensity, age, weight, comorbidities, and so on (Brooks et al. 2014). Adding to this, patients with "complex" attributes are often underrepresented in RCTs, raising clinical uncertainty and the likelihood that patient preferences—including willingness to suffer side effects and to pay for medications—may influence the prescribing choice (Brooks et al. 2014).

Physicians' decisions about which drugs to prescribe are further affected by the evolution of clinical information as new trials are completed—particularly head-to-head studies aimed at establishing the relative efficacy of one drug therapy over another.[21] These ongoing changes in clinical evidence, combined with an expanding set of available products and the accompanying evolution in prices and insurance coverage (e.g., Duggan and Scott Morton 2010), suggest that physicians may turn to drug references that help to ensure patient-specific prescription decisions are based on accurate information.

[21] For example, an RCT completed in 2004 demonstrated that for patients with severe cholesterol abnormality, the incrementally larger reductions achieved by Lipitor resulted in fewer deaths and major coronary events relative to patients taking Pravachol (Cannon et al. 2004). Another such study released in 2008 found that, while Vytorin achieved larger cholesterol reductions than simvastatin, the two drugs were observably identical when it came to the thickness of arterial plaque buildup (atherosclerosis); adding to this, a second study in 2008 found a positive association between Vytorin and cancer (Rossebø 2008) that was later reversed (Cannon et al. 2015).

B. *Prescriptions by US Physicians*

To measure physicians' prescribing of new and existing therapies aimed at cholesterol and lipid control over time, we use physician-level prescription data for the 18 drugs described above from the IMS Health Xponent database.[22] These data are provided at a monthly frequency by drug during the period January 2000 through December 2010, and cover each of the 280,622 US physicians associated with at least ten cholesterol-drug prescriptions during January to December 2010. This low threshold for inclusion implies that our dataset captures essentially the universe of US cholesterol drug prescriptions during this period. For each product and month, we observe the number of prescriptions written by each physician and filled through a US pharmacy. Beginning in January 2006, the data also include information on the method of payment used to fill each prescription (Medicaid, Medicare Part D, cash, or commercial third-party insurance). Importantly, each physician in the dataset is identified by a unique medical education number, name (first name, last name, middle name), and location (a five-digit US zip code). These identifiers enable us to match individual prescribers with their observed pharmaceutical information technology use.

To ensure that our sample includes only those physicians actively prescribing cholesterol drugs during the entire sample period, we restrict attention to the 128,043 physicians that prescribe ten or more statins both during January to December 2000, and during January to December 2010; this allows us to abstract from potential differences in prescribing that may surround a physician's entry into or exit from medical practice, and also ensures that we have adequate data on database adopters' pre-adoption and post-adoption prescribing patterns. The final prescription dataset includes over 200 million observations (132 months $\times$ 128,043 physicians $\times$ up to 18 drugs). Summary statistics appear in Table 2, and additional details regarding data assembly and the Xponent database appear in online Appendix A.2.

C. *Drug Information Access by US Physicians*

To construct an index for the extent of physicians' pharmaceutical information access, we use novel physician-level data from the private firm that owns and operates a prominent electronic reference for pharmaceutical products. The data include a monthly indicator for whether a US physician is a registered user of the reference database, and this suggests the database is widely used: by December 2010, 45.1 percent of sample physicians had established an individual database account (Table 2). The data also include information about registered physicians' actual use of the reference during the sample period; we observe a lower bound on the number of times a physician looks up a cholesterol drug using the database. This proxy is 3.83 on average, and the data indicate that, while 24.2 percent of physicians are registered database users in the average month, only 13.1 percent of physicians use the database to look up one of the cholesterol drugs considered in our study. It is

[22] IMS Health (IQVIA) curates data on additional drug classes; these additional data are not available for the current study due to the unusually large size of the customized data extracts involved.

Table 2—Regression Summary Statistics

Variable	Mean	SD	Min	Max
Physician-Drug-Month level				
Number of prescriptions	4.429	12.721	0	700
Indicator for positive prescriptions	0.355	0.479	0	1
Physician-Month level				
Drug database indicator	0.248	0.432	0	1
Drug database and use indicator	0.133	0.340	0	1
Drug database other adoption share in $zip\ code_{t-1}$	0.131	0.132	0	1
Proxy for intensity of database use	3.829	11.04	0	1,268
Number of unique drugs prescribed	5.304	2.775	1	16
Prescription Herfindahl-Hirschman index (HHI)	0.438	0.223	0.097	1
Generic prescription share	0.043	0.266	0	1
Prescription volume	65.79	66.31	1	2,503
Physician-Drug level				
Months to first prescription	19.12	21.95	0	122
First prescription within initial year indicator	0.352	0.478	0	1
Drug-Month level				
Indicator for new drug, 24 months	0.155	0.363	0	1
General				
Number of physicians	128,043			
Number of drugs, January 2000	6			
Number of drugs, January 2000–December 2010	18			

Notes: This table summarizes the data on physician-level prescriptions and database access used in the analysis. Statistics correspond to US physicians that prescribe a minimum of ten statin or lipid-lowering products both during January–December 2000 and January–December 2010 and that work in a zip code hosting three or more prescribing physicians. The "Drug database indicator" varies by physician-month and is equal to one for physicians that are registered users of the drug database; "Drug database and use" indicates that a physician both has database access and is observed using it to search for information about at least one of the 18 cholesterol drugs during the sample period. "Drug database other adoption share in zip code" varies by physician-month and is the fraction, in the previous month, of other physicians practicing in the same zip code for which "Drug database and use" is equal to one. The intensity of use proxy is a lower-bound on the number of physician-specific database queries corresponding to the cholesterol drugs considered in this analysis. Prescription diversity by physician-month is summarized by the number of unique drugs prescribed and the corresponding Herfindahl-Hirschman index. "First prescription within initial year indicator" takes a value of one for doctors that prescribe the new drug within its initial market year and is otherwise zero. Drugs are considered new if within 24 months of market approval by the US Food and Drug Administration.

Source: Prescription variables are from IMS Health (IQVIA) and database registration data are from a leading US point-of-care medical applications firm.

for this latter group of physicians that database access is likely to be relevant to the cholesterol-drug prescribing outcomes we consider. In Section IV, we thus consider whether the observed intensity of database use explains variation in its impact on prescribing.

The drug reference we study contains information that is, in principle, relevant for improving the match between patient characteristics and available pharmaceutical products. At any point in time, the drug reference contains detailed information about each available US FDA-approved medication. This information is obtained from the medical literature, specialist recommendations, clinical guidelines, manufacturer labeling, standard medical references, and FDA drug safety alerts and is updated continually. The results of this ongoing research are condensed into drug-specific monographs that may be accessed through the electronic database

interface. Beyond standard clinical information such as contraindications, cautions, adverse reactions, safety, monitoring, and pharmacology, the reference monographs also include a set of additional variables for each product that may affect prescribing decisions. Specifically, the monographs include retail pricing and formulary status information for each drug, drug interaction information, FDA warnings, and off-label and pediatric usage guidelines. Each physician customizes the tool with respect to formularies, selecting those relevant to their decision needs; it is then straightforward, for any drug, to check copay tiers, formulary alternatives, generic substitutions, criteria for prior authorizations, and quantity limits—facets of a formulary that are subtle but often have significant consequences for patient costs. The database includes separate entries for each branded product and each generic product based on product-specific information such as available formulations, dosing, indications, manufacturer, and pricing. The database is updated to reflect both the current set of products and formulary details, as well as the current state of knowledge regarding drug characteristics and clinical practice. Importantly, information for new drugs becomes available around the time the drug is released by the FDA for commercial prescription.

Because the drug reference combines available information into a single, current monograph rather than contributing new or proprietary drug information, it is best viewed as a tool that makes it convenient for physicians to quickly access condensed clinical, insurance, and pricing information about a drug. Doctors commonly use the reference to check dosages, contraindications, and coverage details, but rely on other sources, such as medical journals or more encyclopedic references, for information such as a drug's results in clinical trials.

The database provides certain forms of clinical guidance to prescribers. While it recommends statins as a first-line treatment for use in reducing LDL-C (low-density lipoprotein cholesterol), it is important to note that guidance in terms of which among the available drugs to prescribe is highly patient-specific. For example, the database identifies each statin as falling into one of three groups: high-intensity, moderate-intensity, and low-intensity. These groups are defined based on numeric LDL-C reduction targets that are specific to a patient. The reference indicates that a patient in the high-intensity category (daily dosage lowers LDL-C by over 50 percent on average) should receive either atorvastatin with a dosage of 40–80 mg/day (Lipitor) or rosuvastatin with a dosage of 20–40mg/day (Crestor).[23]

The electronic drug reference database we study is a stand-alone technology, not embedded into other health information technology systems, and is accessed almost exclusively by an individual (rather than institutional) subscription. This is important because if the database were part of a broader health information technology platform—for example one that includes features that facilitate patient and insurer billing (see Agha 2014)—it would be difficult to know whether estimates corresponding to database use indicate effects of drug information access or effects

[23] Notice that evaluating whether doctors with database access better adhere to these guidelines than those without access would require observing (i) patient-specific LDL-C reduction goals for each doctor-month, and ideally also (ii) the dosage for each dispensed prescription. Because we are unable to observe either of these details, it is not possible in our analysis below to assess directly whether adherence to prescribing guidelines improves with database usage.

of some other technology that is part of the same health IT system that the doctor accesses.[24]

For our study, it is critical to understand what drives database adoption. Figure 1 indicates that use of the reference database during the sample period is not random, but differs according to observable doctor characteristics.[25] Throughout the sample period, physicians are more likely to have adopted the database if they had graduated from a top-ranked US medical school (panel A) and had graduated recently (panel B); males are also more likely to adopt (panel C). Doctors in obstetrics and gynecology (panel E) and those practicing in the US South (panel F) appear systematically slower to adopt the database.

Presently, new physician adopters tend to learn about the reference while in medical school. However, because the physicians in our dataset had all completed medical school before the database became available, their adoption decisions are more likely to have been driven by marketing or peer effects. Documents filed along with the reference firm's initial public offering state that its marketing strategy was, in fact, an informal "word-of-mouth" approach, and that throughout the sample period, the network of reference users grew over time primarily through users telling friends and colleagues about its value. The filings state that this strategy had been both highly effective and inexpensive relative to the alternative of hiring a dedicated sales force.[26] Thus, while our data indicate that physicians are visibly idiosyncratic in their adoption timing, a doctor is much more likely a user if a high share of other doctors in her zip code are also users; this is consistent with the firm's reported marketing strategy. Moreover, only 16 percent of the variation in the time to adoption is explained by zip code fixed effects, indicating that within-zip-code dynamics are quite important.

Regarding adoption motives, the main reason doctors cite for registering is convenience: database use reportedly yields meaningful time savings. By contrast, it is unlikely that price was an important factor in physicians' adoption decisions. Access was always available through a free version of the database application, which included the core drug reference tools (e.g., dosage lookups) that are relevant to our study. Additional features were available with a paid subscription, but the annual fee for this version was low (never above $200).

It should also be noted that database adoption appears to be mostly an individual decision, even for doctors in group practices. Large clinics and physician groups sometimes purchase site licenses for institution-level access to the database as part of broader IT initiatives; however, some of the benefits of using the database require individual registration, and most doctors therefore have individual accounts even if their group or clinic has a site license. Nevertheless, in order to check whether

[24] Relatedly, other drug references exist, including the Micromedex and UpToDate Lexicomp databases. These are imperfect substitutes for what we consider: these databases do not contain drug price or formulary-specific coverage information and are typically accessed by institutional subscription. Moreover, our data do not include information on doctors' potential usage of these alternative platforms. Thus, as a matter of interpretation for our estimation results below, if a doctor tends to begin using all databases around the same time, our estimates would reflect the influence of information from all platforms rather than just the one we observe directly.

[25] Physician characteristics were obtained from the Centers for Medicare and Medicaid Services Physician Compare database and were matched based on physician first name, last name, and five-digit zip code.

[26] By contrast, sales force marketing is standard for new pharmaceuticals. See Datta and Dave (2017).

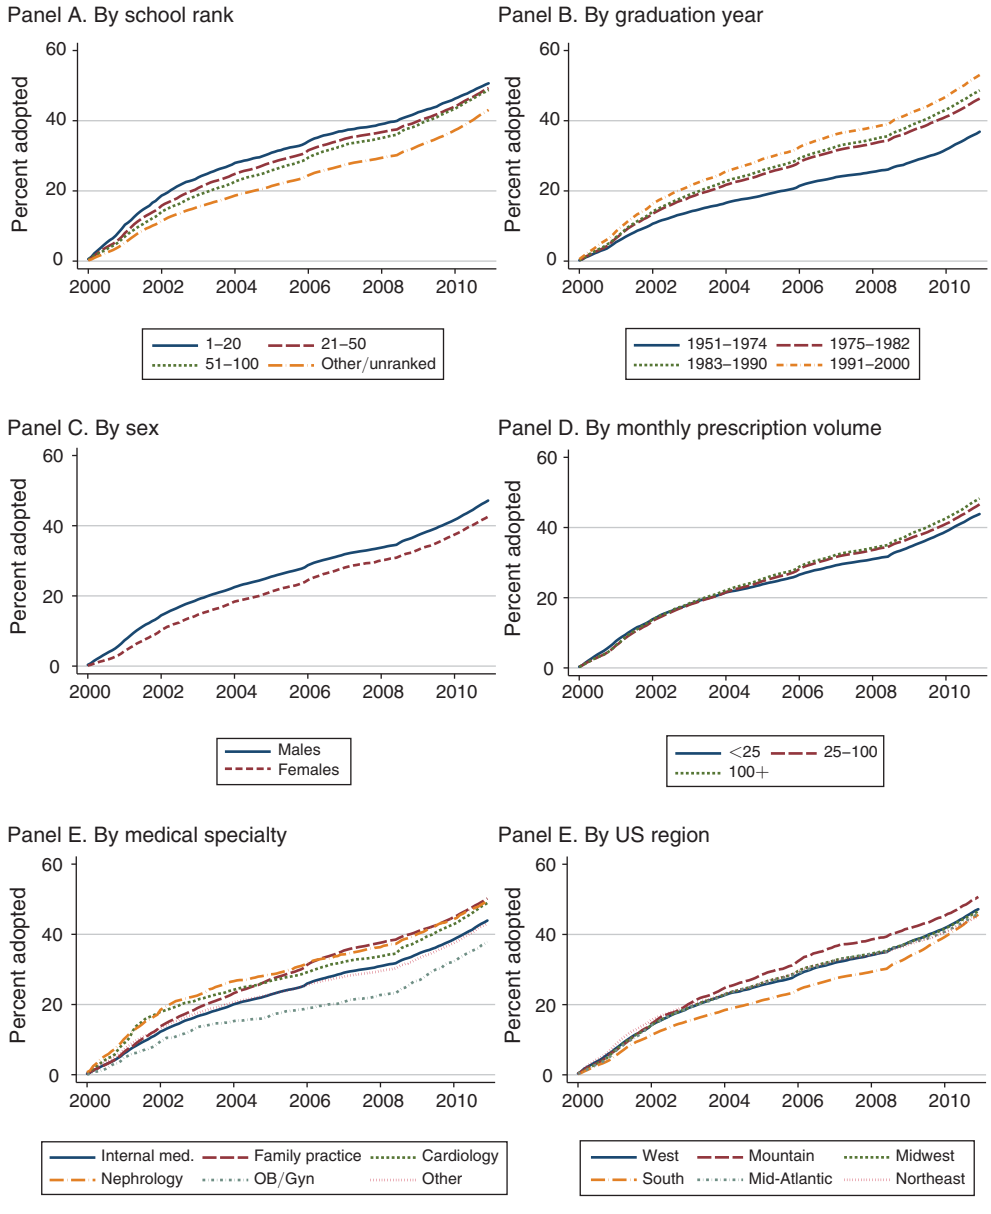

FIGURE 1. DRUG DATABASE DIFFUSION CURVES, US PHYSICIANS, JANUARY 2000–DECEMBER 2010

Notes: This figure plots the fraction of the approximately 67,000 sample US physicians included in the CMS Physician Compare database that are also registered users of the electronic drug reference database by the date indicated, and shows the extent to which adoption rates differ across physicians according to their observable characteristics.

Sources: Database registration data are from the drug reference database firm. Medical school rank is determined based on data from the *US News and World Report* service, and all other variables are from the CMS Physician Compare database.

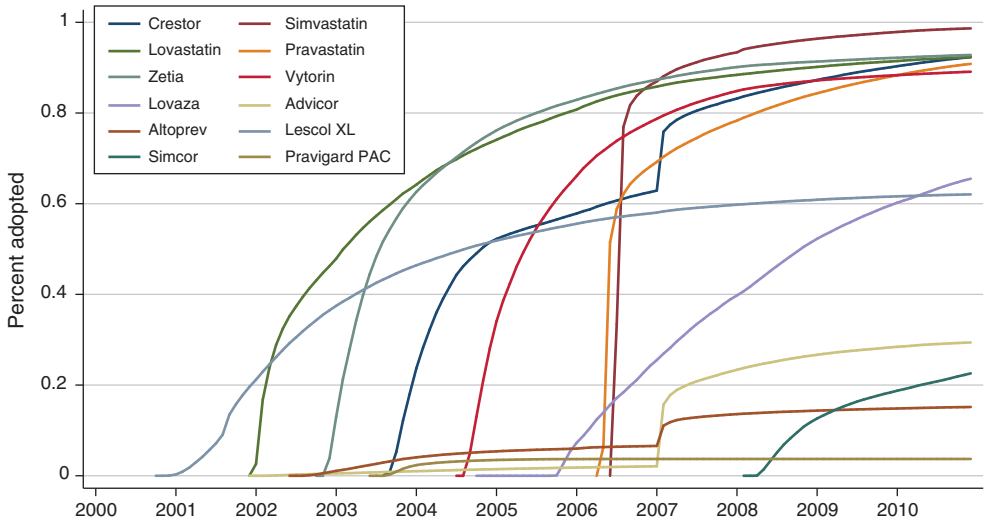

Figure 2. Drug Diffusion Curves by Drug, US Physicians, January 2000–December 2010

Notes: This figure plots the fraction of the 128,043 sample US physicians that are associated with at least one prescription of the new drug indicated by the date marked on the horizontal axis, and shows the extent to which adoption rates differ across products. Market approval dates by drug are listed in Table 1. The prescription data cover January 2000 through December 2010 at a monthly frequency and are from IMS Health (IQVIA).

doctors practicing in groups tend to synchronize their database adoption—which would suggest the influence of a group-level adoption decision—we used the 2014 CMS Physician Compare database to identify doctors who were likely working in the same practice during our sample period. Among over 7,000 groups we identified, just 38 were ones in which all doctors in the group adopted the database at the same time. In light of this, it seems unlikely that site-level access or group adoption decisions are primary drivers of the physician-level database use we consider.

D. *Descriptive Evidence*

The data provide suggestive indications that incomplete information may affect physicians' prescribing as well as the rate and extent of new product diffusion, depicted in Figure 2. Consider the statistics presented in Table 3, which quantify differences in prescribing across US physicians for the class of cholesterol medications evaluated. The statistics in panel A provide evidence for the December 2010 cross section. It is apparent that the pronounced variation in cholesterol-drug prescribing previously found among Medicare patients (e.g., Munson et al. 2013, Brooks et al. 2014) is also present within the overall population both across zip codes (columns 5–8) and individual physicians (columns 1–4). The share of prescriptions accounted for by generics ranges from zero to one in column 4; moreover, while the average physician prescribes a generic in 56.4 percent of cases, the standard deviation is also large (24.3 percent) and spans the full range from zero (fifth percentile physician) to 100 percent (ninety-fifth percentile). The relative heterogeneity across doctors is

Table 3—Descriptive Statistics

	Physician level				Zip code level			
Product	lovastatin	pravastatin	simvastatin	generic	lovastatin	pravastatin	simvastatin	generic
Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A	Share in total Rx, by physician				Share in total Rx, by zip code			
Final month, December 2010								
Mean	0.059	0.091	0.414	0.564	0.065	0.098	0.427	0.591
SD	0.116	0.135	0.238	0.243	0.083	0.094	0.159	0.161
Fifth percentile	0	0	0	0	0	0	0.180	0.336
Twenty-fifth percentile	0	0	0.258	0.427	0.017	0.041	0.335	0.502
Median	0.016	0.047	0.410	0.581	0.041	0.078	0.423	0.596
Seventy-fifth percentile	0.066	0.121	0.551	0.719	0.084	0.129	0.515	0.687
Ninety-fifth percentile	0.271	0.338	0.858	1	0.209	0.259	0.271	0.834
Panel B	Generic Rx share, by physician-molecule				Generic Rx share, by zip-code-molecule			
Six months after generic release								
Molecule-specfic branded drug	Mevacor	Pravachol	Zocor		Mevacor	Pravachol	Zocor	
Mean	0.828	0.820	0.862		0.822	0.827	0.870	
SD	0.338	0.279	0.208		0.192	0.197	0.150	
Fifth percentile	0	0	0.448		0.491	0.490	0.582	
Twenty-fifth percentile	1	0.714	0.800		0.737	0.756	0.819	
Median	1	1	0.949		0.861	0.866	0.914	
Seventy-fifth percentile	1	1	1		1	1	0.977	
Ninety-fifth percentile	1	1	1		1	1	1	
Panel C								
Final month, December 2010								
Molecule-specific branded drug	Mevacor	Pravachol	Zocor		Mevacor	Pravachol	Zocor	
Mean	1.000	0.993	0.997		1.000	0.995	0.998	
SD	0.019	0.059	0.028		0.005	0.030	0.011	
Fifth percentile	1	1	0.995		1	0.976	0.990	
Twenty-fifth percentile and above	1	1	1		1	1	1	

Notes: This table describes prescription heterogeneity across US physicians and the US zip codes they occupy. Panel A describes prescribing in December 2010 across all physicians (columns 1–4), and across US zip codes (columns 5–8). Panels B and C describe physicians' within-molecule substitution toward generics for lovastatin (column 1), pravastatin (column 2), and simvastatin (column 3); columns 5, 6, and 7 provide analogous statistics by US zip code. Panel B describes this substitution six months after the generic release in question, while panel C describes prescribing in the final sample period, December 2010. The upper-left number in panel A (mean, lovastatin, 0.059) is the average, across physicians, in the fraction of cholesterol drug prescriptions in December 2010 that are accounted for by generic lovastatin; the upper-left number in panel B (mean, lovastatin, 0.828) is the average, across physicians, in the fraction of Mevacor plus generic lovastatin prescriptions that are accounted for by generic lovastatin in October 2002, six months after expiration of the Mevacor patent; the upper-left number in panel C is the analogous statistic for December 2010.

Sources: Generic approval dates are from the US Food and Drug Administration; all other variables are from IMS Health (IQVIA).

even wider for the specific drugs described in columns 1–3; and across all columns, the idiosyncratic behavior of individual physicians appears to be important: for the vast majority of zip codes, there is substantial within-zip-code variation in generic shares across local physicians (Table A.1 in online Appendix).

Even if physicians were perfectly informed, variation in prescribing could result from an uneven distribution of patient characteristics. For example, Lipitor is a high-intensity statin that may be preferable for patients with a severe cholesterol abnormality, the incidence of which may cluster geographically. Similarly, risk-averse patients may prefer an established drug over a new one—even if the new drug is simply a new generic version—if they perceive the quality of a new product as uncertain relative to another option. Such underlying patient heterogeneity may partially explain the slow and incomplete diffusion of new drug varieties, which is

apparent for each new drug except simvastatin (Figure 2 and Table 1, column 3). It may also explain why a substantial fraction of the variation across doctors observed in columns 1–4 remains even after aggregating to the zip code level (columns 5–8).

Unobserved patient heterogeneity likely explains some of this variation in prescribing, but columns 1, 2, and 3 indicate that additional factors are also likely present. Specifically, these columns assess within-physician changes in the prescription of new generic products. The advantage of focusing on these columns is that it is possible to compare prescribing of a branded product with its molecularly equivalent generic—two distinct drugs that have no relevant clinical differences. By examining changes over time in the generic share of molecule-specific prescriptions (e.g., the share of simvastatin plus Zocor prescriptions that are accounted for by generic simvastatin), it is possible to determine whether stable patient heterogeneity is likely to be the only explanation for variations in care. For each of the three generic drug introductions (lovastatin, pravastatin, and simvastatin), the data indicate that physicians differ in their use of generics in the short run, six months after generic entry, and that substitution toward generics is initially incomplete at this point (panel B). By contrast, in the long run, physicians differ substantially less: nearly complete substitution toward generics is observed for each of the three products (panel C).[27] This pattern of *delayed substitution* between two molecularly equivalent products strongly suggests factors other than time-invariant patient heterogeneity contribute to prescribing differences among cholesterol drugs and is consistent with the influence of information frictions.

Beyond cost implications, these same factors may impede the diffusion of new non-generic therapies, with consequences for health outcomes. The data indicate that physicians are slow to begin prescribing new molecular entities, new drug combinations, and new dosage forms—branded products not facing generic competition. Figure 2 shows that diffusion curves differ considerably across new drugs; Figure 3 plots the gradual diffusion of Crestor across US zip codes; and Table 1 describes how the time lag in months between a drug's approval and its initial prescription varies across US physicians for each drug introduction. The average physician delays prescribing a new drug for 20.3 months among the new products considered in our analysis. The standard deviation is even larger (22.1 months), and this adoption lag ranges from zero to 122 months, indicating some physicians adopt immediately and others had yet to adopt the first new drug by the end of our sample period (Table 2).

Unlike the Dartmouth Atlas and Cooper et al. (2015) studies, which cover a comprehensive set of treatments and analyze regional differences in the cost and quality of care, we analyze behavior at the physician level and focus on the specific clinical decision of which cholesterol drug to prescribe. However, it is nevertheless useful to ask whether the patterns we observe for this context are consistent with the broader treatment patterns documented by the Atlas. Aggregating our physician-level data to Hospital Referral Regions (HRRs) and comparing against data from the Dartmouth Atlas project, we find that locations with high generic prescription

[27] By December 2010, physicians had broadly switched away from prescribing Mevacor, Pravachol, and Zocor. However, six months after each respective patent expired, generic prescribing was far less prevalent for each molecule, though the generic version was in each case already substantially less expensive.

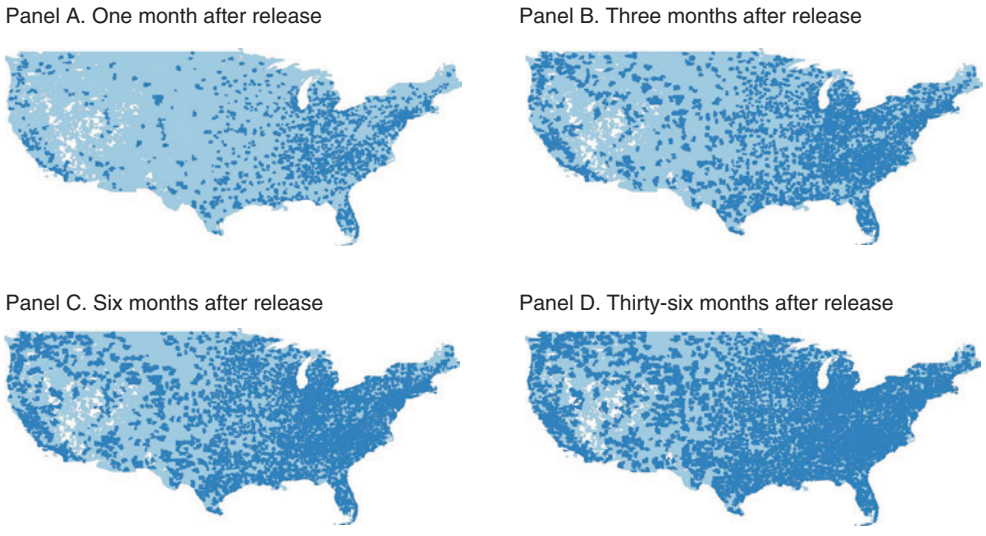

FIGURE 3. HETEROGENEITY IN THE INITIAL USE OF A NEW MEDICAL TECHNOLOGY BY US ZIP CODE

Notes: This figure illustrates the gradual diffusion of a new pharmaceutical drug, the statin Crestor, across zip codes within the continental United States. Dark shades indicate zip codes in which at least one prescription of Crestor has been written and filled, light shades indicate zip codes in which Crestor has not yet been prescribed; areas shaded white contain no data. The four panels correspond to four points in time following the initial market introduction of Crestor in August 2003. These four points are September 2003 (panel A), November 2003 (panel B), February 2004 (panel C), and August 2006 (panel D).

Source: Prescription data are from IMS Health (IQVIA).

shares in our dataset also have high generic shares for Medicare overall (correlation 44.4 percent), as well as lower per capita medical spending for prescription drugs (correlation 23.9 percent) and non-drug healthcare (correlation 12.4 percent). As noted above, these patterns underscore the highly systematic nature of US disparities in care and suggest that the disparities in physicians' prescribing of cholesterol drugs may reflect some of the same factors that drive disparities in treatment decisions more broadly.

II. Empirical Strategy

In this section we provide a conceptual framework indicating how physician information and prescribing outcomes may be related. We describe the model implications and restrictions that guide our approach to estimating the treatment impact of physicians' database access.

A. *Conceptual Framework*

Consider a baseline model in which physician i faces a period-t choice over which drug to prescribe for each of her patients $n = 1, 2, \ldots, N_{it}$. Like other economic studies of prescribing decisions, suppose that physician i makes this decision for each patient by selecting the single drug $j \in \{1, 2, \ldots, J_t\}$ available at t

that maximizes patient utility according to physician-i information.[28] Specifically, suppose that the true utility derived by patient n from drug j at t is $u_{njt} \equiv \theta_{jt} + V_{njt}$, which combines the quality of drug j that is both known at t and common across patients (θ_{jt}) with the quality of j that is unknown and partially specific to patient n (V_{njt}). The first of these terms (θ_{jt}) thus captures the accepted wisdom at t about the efficacy, costs, side effects, and so on of drug j for the average patient, while the second reflects novel information that may, in part, be relevant to the match between j and patient n. In particular, suppose that V_{njt} combines two terms: $V_{njt} \equiv v_{jt} + \epsilon_{njt}$, where v_{jt} is a drug-specific value—a revision to accepted wisdom about the quality of drug j—and where ϵ_{njt} reflects the quality of the match between patient n and drug j. We assume the physician is only partially informed about V_{njt}, to a degree indexed by a parameter ϕ_{it}; she bases her prescribing decision on a partial observation of u_{njt} given by

$$\hat{u}_{njt} \equiv \theta_{jt} + \left(1 - e^{-\phi_{it}}\right) V_{njt} = \theta_{jt} + \left(1 - e^{-\phi_{it}}\right)\left(v_{jt} + \epsilon_{njt}\right). \tag{1}$$

Physicians with a higher value of ϕ_{it} in equation (1) are more responsive to information about drug quality that is not commonly known at t (v_{jt}), and about the patient-specific match (ϵ_{njt}). In particular, (1) implies physicians with no special information ($\phi_{it} = 0$) are insensitive to V_{njt} and thus prescribe the same drug—that with the highest θ_{jt}—for all patients, while physicians who are fully informed ($\phi_{it} \to \infty$) respond to V_{njt} perfectly.

If we assume that the ϵ_{njt} follow an i.i.d. Type-1 Extreme Value distribution, it is straightforward in this simple setup to show that the probability physician i prescribes drug j for patient n at t depends on the information index ϕ_{it} as follows:

$$p_{jt}(\phi_{it}) = \frac{\exp\left\{\frac{\theta_{jt}}{1 - e^{-\phi_{it}}} + v_{jt}\right\}}{\sum_{k=1}^{J_t} \exp\left\{\frac{\theta_{kt}}{1 - e^{-\phi_{it}}} + v_{kt}\right\}},$$

and that, accordingly, the probability P_{ijt} that drug j is prescribed by physician i at least once during period t is

$$P_{ijt}(\phi_{it}) \equiv P\{X_{ijt} > 0\} = 1 - P\{X_{ijt} = 0\} = 1 - \left(1 - p_{jt}(\phi_{it})\right)^{N_{it}}, \tag{2}$$

where X_{ijt} is the number of physician-i prescriptions written for drug j during period t.[29] Moreover, starting from the introduction date t_j^0 of a new drug j, the

[28] See, for example, Dickstein (2018), Crawford and Shum (2005). Unlike these papers, we do not observe patient-level information; this precludes estimating a model of learning within each patient-physician pair.

[29] Qualitatively identical results hold under more general assumptions regarding the distribution of ϵ_{njt}; the Type-1 Extreme Value assumption is thus imposed here only for expositional simplicity. A realistic alternative would be to allow for persistence in the ϵ_{njt} draws, reflecting that the chronic nature of cholesterol and lipid disorders implies physicians often treat the same patient for multiple successive periods. In our analysis to follow, we thus consider the possibility that the prescription outcomes we evaluate are persistent.

expected number of periods T_{ij} that lapse before drug j is prescribed at least once by physician i is

$$(3) \qquad E[T_{ij}] = \sum_{t=t_j^0}^{\infty} (t - t_j^0) P_{ijt}(\phi_{it}) \prod_{s=t_j^0}^{t-1} (1 - P_{ijs}(\phi_{is}))$$

$$= \sum_{t=t_j^0}^{\infty} (t - t_j^0) \left(1 - \left(1 - p_{jt}(\phi_{it})\right)^{N_{it}}\right) \prod_{s=t_j^0}^{t-1} \left(1 - p_{js}(\phi_{is})\right)^{N_{is}},$$

which also depends on ϕ_{it}, as does the expected number of unique drugs M_{it} prescribed by physician i during t:

$$(4) \quad E_t[M_{it}] \equiv E_t\left[\sum_{j=1}^{J_t} \mathbf{1}\{X_{ijt} > 0\}\right] = \sum_{j=1}^{J_t} P_{ijt}(\phi_{it}) = \sum_{j=1}^{J_t} \left(1 - \left(1 - p_{jt}(\phi_{it})\right)^{N_{it}}\right).$$

Within this framework, we regard the electronic database as a technology that increases a physician's ϕ_{it}, which is otherwise unobserved. The database is continuously updated, so users of the database are more likely aware of any new information about the drug, including price changes, new warnings, or new results about its efficacy for different patient types. The database also allows doctors to look up a drug's current formulary status for a specific patient's insurance plan, so database users should also be more responsive to differences in, and changes in, match quality across patient-drug pairs. From equation (2), if database use indeed increases ϕ_{it}, it impacts the probability drug j is prescribed: whether P_{ijt} increases or decreases for drug j depends on the distribution of V_{njt} across drugs j and patients n. In general, P_{ijt} will increase for drugs with high values of v_{jt} relative to other drugs; alternatively, if all $v_{jt} = 0$, an increase in ϕ_{it} raises P_{ijt} for all drugs (due to the ϵ_{njt}) except for that with the highest θ_{jt}. Similarly, equation (3) implies the expected number of periods that pass before drug j is prescribed declines in ϕ_{it} whenever P_{ijt} increases in ϕ_{it}. That is, if a permanent increase in ϕ_{it} causes a permanent increase in P_{ijt} for drug j, then it also causes a decrease in T_{ij}. The impact of an increase in ϕ_{it} on the number of distinct drugs prescribed depends on the sum of derivatives $\partial P_{ijt}(\phi_{it})/\partial \phi_{it}$ across drugs j in equation (4). If a higher ϕ_{it} implies increased sensitivity to patient-specific match quality ϵ_{njt}, for example, P_{ijt} would increase for most drugs j, and diversity of prescribing would also rise.[30]

It is important to note that doctors who regularly prescribe cholesterol medications will be aware of most drugs' clinical attributes. However, if patient-specific economic details such as the pricing and formulary status of a drug evolve substantially over time, or if news about negative drug interactions and other adverse reactions emerges only gradually, doctors may prefer to look up these drug attributes prior to writing a prescription. For newer, less familiar drugs, doctors may also be

[30] See also Berndt et al. (2015).

inclined to look up details like dosage, and it is for these drugs that one may expect differences between u_{njt} and $\hat{u}_{njt}$ to be particularly relevant.

B. *Econometric Model*

One natural approach to evaluating the influence of increases in ϕ_{it} due to database adoption would be to directly estimate equations derived from the conceptual model above. However, the model indicates it is important to control not only for the number of drugs J_t and i's prescribing intensity N_{it}, but also for unobserved drug quality θ_{jt} and unobserved determinants of ϕ_{it}. Given the size of the dataset, handling the nonlinearity implied by (2) in the presence of multiple sets of fixed effects is computationally infeasible. We therefore estimate the effects of database use through specifications that are guided by the model, but linear.

In particular, we consider three main linear specifications corresponding to each of the three observable outcomes discussed above: the new drug adoption lag T_{ij}, prescription diversity M_{it}, and prescription probabilities P_{ijt}. Our estimation approach does not impose the restrictions that link P_{ijt} with T_{ij} and M_{it} in the model; as a result, comparing our estimates across these outcomes is qualitatively informative regarding the fit of the model.

We first assess the time lapse T_{ij} between the initial market release of drug j and its first prescription by physician i as in Coleman, Katz, and Menzel (1957) using the following equation:

$$P\{T_{ij} \leq 12\} = \eta_i + \eta_{zj} + \beta Z_{ij} + \delta N_{it(j)-1} + \varepsilon_{ij}, \tag{5}$$

where T_{ij} is measured in months, $P\{T_{ij} \leq 12\}$ is the probability that j is prescribed by i within twelve months of release, and where η_{zj} and η_i are zip-code-drug and physician fixed effects, respectively.[31] The variable Z_{ij} indicates whether doctor i has database access at the time drug j is first introduced, and $N_{it(j)-1}$ is i's total prescription volume for cholesterol drugs in the month preceding j's introduction.

Both (3) and (5) are expressed at the doctor-drug level and, but aside from the functional-form differences mentioned above, the two equations are connected. The estimating equation essentially considers ϕ_{it} to be a function of database use (Z_{it}), physician fixed effects (η_i), and zip-code-month fixed effects (η_{zt}) reflecting local changes in access to information. We arrive at equation (5) by noting that the evolution of time is, in doctor-drug space, marked by the sequential introduction of each new drug j, and that the η_{zj} we therefore include take the place of η_{zt} while also accounting for the drug quality effects θ_{jt} in (3). In addition, like (3), (5) includes i's prescription volume $N_{it(j)-1}$, and J_t is absorbed by the η_{zj}.

Equation (5) is estimated on the subset of drugs first introduced during the sample period. Finding that the coefficient of interest β is positive would indicate that when a physician obtains database access, she significantly increases her likelihood

[31] With T_{ij} as the dependent variable, it is necessary to address truncation, which is more pronounced for drugs introduced late in the sample period. To apply a uniform truncation rule, a significant number of observations must be omitted; hence T_{ij}, though more direct, is not our preferred dependent variable.

of prescribing a new drug within its first year, relative to before access began. This would be consistent with database use increasing ϕ_{it} and the probability P_{ijt} of prescribing a new drug j at t: $\partial P_{ijt}(\phi_{it})/\partial\phi_{it} > 0$. Notice that the inclusion of physician fixed effects implies that β is identified using within-doctor variation over time; i may be a database user at the time drug j is first introduced, but may not yet be a user upon the introduction of drug j'. These effects are important if stable, unobserved physician characteristics determine both physician-specific database use Z_{ij} and the rate of drug adoption T_{ij} (e.g., early adopters). Possible common unobserved random shocks that are local and correlated with database adoption are further accounted for by clustering standard errors at the zip code level.

Nevertheless, even with these fixed effects and clusters, there could be time-varying factors such as local technology adoption propensities (Skinner and Staiger 2007) or pharmaceutical advertising that jointly determine, or are correlated with, both physician i's database use and her rate of new drug adoption. The η_{zj} in (5) partially address this by accounting for differences across zip-code-drug pairs in doctors' average first-prescription timing, which in this context would be correlated with Z_{ij}. However, if the omitted factor is idiosyncratic across physicians, even within a zip code, then $\text{cov}(Z_{ij}, \varepsilon_{ij}) \neq 0$ and (5) will fail to yield a consistent estimate of β. We return to this in describing our instrumental-variables estimates in Section IV.

Second, building from (4), we consider the possibility that information access could affect physician i's knowledge about the match quality between drug j and patient n, inducing better-informed physicians to prescribe a more diverse set of products than less-informed peers. To assess this possibility, we determine the number of unique drug products $M_{it} \equiv \sum_{j\in\mathcal{J}_t} \mathbf{1}\{X_{ijt} > 0\}$, where $\mathbf{1}\{X_{ijt} > 0\}$ is an indicator for whether physician i writes at least one prescription for drug j during month t, and evaluate the following specification:

$$M_{it} = \eta_i + \gamma_i \times t + \eta_{zt} + \beta Z_{it} + \delta N_{it-1} + \varepsilon_{it}, \tag{6}$$

where $\gamma_i \times t$ is a doctor-specific time trend and all variables are as defined above. The value of M_{it} is low when the prescriptions of physician i are concentrated within a narrow subset of products during month t, and is high when prescribing is diverse. Finding that β is positive in equation (6) would thus indicate that database access is associated with higher product diversity among physician i's prescriptions—this would further be consistent with an increase in the overall *quality* of doctor-i prescribing whenever the common component ν_{jt} in (1) is small relative to ϵ_{njt}. Notice also that, through (4), this occurs only when the period-t prescription probability P_{ijt} rises more, on average, than it falls—that is, when $\sum_{j=1}^{J_t} \partial P_{ijt}(\phi_{it})/\partial\phi_{it} > 0$. Setting aside functional forms, the connection between (6) and (4) again rests on the idea that the information index ϕ_{it} is a function of database use, physician fixed effects, and zip-code-month fixed effects. We control for N_{it-1} directly, and for J_t through η_{zt}. Standard errors are clustered by zip code to allow for local unobserved shocks correlated with Z_{it}.

We also estimate (6) replacing M_{it} with the Herfindahl-Hirschman index (HHI_{it}) as an alternative dependent variable; an advantage of this alternative is that it

simultaneously captures both intensive- and extensive-margin effects of information on prescribing. Notice that β in (6) is again identified using within-physician variation over time in information access Z_{it}. The zip-code-month fixed effects further help to account for changes over time in unobserved, location-specific determinants of prescribing diversity; these are particularly important if patient characteristics—such as insurance coverage, mandatory substitution laws, patient preferences, and disease severity—or other local factors evolve in ways that affect prescribing and are correlated with measured physician technology adoption. The $\gamma_i \times t$ trends further account for doctor-specific linear changes over time in these unobserved determinants of prescribing outcomes.

Finally, we evaluate directly the impact of information access on P_{ijt}, the probability that drug j is prescribed by doctor i at t. It is of particular interest to understand how database users' P_{ijt} values across new and old drugs j change after database adoption. Moreover, because new patent-protected products differ from new generics in both cost and novelty, access to the database may also tilt prescribing based on the patent status of a product. We thus evaluate whether physicians using the database are more or less likely to prescribe a product of a given type using the following specification:

$$
\begin{aligned}
(7) \qquad P\{X_{ijt} > 0\} = {} & \eta_{jt} + \eta_i + \gamma_i \times t + \eta_{zt} + \delta N_{it-1} \\
& + \left[\beta_0 Gen_j + \beta_1 (1 - Gen_j)\right] \times Z_{it} \times New_{jt}^{\tau} \\
& + \left[\beta_2 Gen_j + \beta_3 (1 - Gen_j)\right] \times Z_{it} \times \left(1 - New_{jt}^{\tau}\right) + \varepsilon_{ijt},
\end{aligned}
$$

where $P\{X_{ijt} > 0\}$ is the probability that physician i writes at least one prescription for drug j during month t, Gen_j is an indicator that is equal to 1 if product j is a generic variety, and New_{jt}^{τ} indicates whether drug j is within τ months of its initial approval for US sale. The main coefficients of interest β_0, β_1, β_2, and β_3 jointly capture the association between database use Z_{it} and prescribing propensity for both new drugs (β_0, β_1) and established products (β_2, β_3), where finding $\beta_0 > 0$ would indicate that database users are more likely to prescribe a given drug j that is both new (within τ months of initial market release) at t and generic, relative to other physicians. Similarly, finding that $\beta_1 > 0$ would indicate that database users are more likely to prescribe a new, branded product j. Note that the estimates of (7) have implications for T_{ij} and M_{it} through (3) and (4) above.

Equation (7) includes three sets of fixed effects, in line with (2). Drug-month effects η_{jt} account for the average perceived quality of drug j across physicians at t (θ_{jt}), which may depend on factors such as drug potency and side effects known at t, as well as the average expected pharmacy price at t. As with our other estimating equations, we further include physician fixed effects η_i that absorb any individual characteristics affecting ϕ_{it} or the prescribing propensity such as location, patient composition, and physician age, education, and medical specialty. The coefficients of interest β are thus identified primarily from within doctor variation over time in information access Z_{it}. Zip-code-month fixed effects η_{zt} absorb any dynamic, location-specific determinants of ϕ_{it} or prescribing that may be correlated with

TABLE 4—TIME TO FIRST PRESCRIPTION OF A NEW DRUG, US PHYSICIANS, 2000–2010

Dependent variable	Indicator for prescription within first year of drug introduction					
	All physicians				Eventual users	
	(1)	(2)	(3)	(4)	(5)	(6)
$Database_{ij}$	0.0204	0.0082	0.0152	0.0010	−0.0015	−0.0029
	0.0013	0.0021	0.0014	0.0022	0.0022	0.0028
$Database_{ij} \times Generic_j$			0.0183	0.0208	0.0136	0.0131
			0.0022	0.0023	0.0039	0.0040
$Prescription\ volume_{it(j)-1}$	0.0073	0.0033	0.0073	0.0033	0.0066	0.0029
	0.0001	0.0001	0.0001	0.0001	0.0002	0.0001
Physician fixed effects	No	Yes	No	Yes	No	Yes
Zip-code-drug fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,510,985	1,510,985	1,510,985	1,510,985	290,898	290,898
R^2	0.5133	0.6132	0.5134	0.6133	0.5976	0.6771

Notes: This table provides least-squares estimates of equation (5) for US physicians' prescription of 12 cholesterol drugs first approved for US sale during January 2000–December 2008 (Table 1). The binary dependent variable captures the time lapse between FDA approval of drug *j* and physician *i*'s initial prescription of it, taking a value of 1 if initial prescription occurs within a year of FDA approval; specifications are included for the full sample of physicians (columns 1–4) and for the subset of physicians that eventually adopt and use the electronic reference to search for information about cholesterol drugs (columns 5–6). Database is the "Drug database and use indicator" variable described in Table 2, and takes a value of 1 for a physician user with database access at the time drug *j* receives FDA approval. Generic indicates the products pravastatin, lovastatin, and simvastatin. Regressions include zip-code-drug (columns 1–6) and physician (columns 2, 4, 6) fixed effects as well as the cholesterol drug prescription volume for physician *i* in the month prior to drug *j*'s introduction. Results are robust to logistic estimation, and are qualitatively identical when replacing the dependent variable with an indicator for first prescription within two years. Standard errors clustered by zip code appear below each point estimate; results are robust to clustering errors by physician.

physicians' database use, while including $\gamma_i \times t$ allows for doctor-specific trends in prescribing that may be correlated with database use. Standard errors are adjusted for clustering at the zip code level, and we control for N_{it-1} directly.

III. Main Results

A. *Time to First Prescription*

We begin by evaluating the relationship between a physician's database use and whether she adopts a new drug *j* within a year of its release. The model indicates that physician *i* is faster to begin prescribing *j* if she is a database user, for any new drug *j* satisfying $\partial P_{ijt}(\phi_{it}) / \partial \phi_{it} > 0$. For these drugs, we thus expect a positive coefficient on Z_{ij} ($Database_{ij}$), where Z_{ij} takes a value of 1 if physician *i* has access to the drug reference database at the time a new drug *j* receives approval for sale in the US market, and is otherwise zero.[32]

Estimates of (5) appear in Table 4. Columns 1 and 2 support the idea that database users are more likely, on average, to begin prescribing newly-approved drugs early, within their first year. The estimated coefficient on Z_{ij} is positive and highly

[32] Throughout Section IVA, $Database_{ij}$ (Z_{ij}) is the "Drug database and use indicator" described in Table 2.

significant in both columns, suggesting users are 2.04 percentage points more likely than nonusers to write their initial prescription for a new drug within its first year (column 1); the estimate changes to 0.82 percentage points if we include physician fixed effects (column 2).[33] The data also confirm that doctors with large prescription volumes N_{it-1} are also significantly faster to begin prescribing a new drug, consistent with the model's qualitative predictions.

Columns 3 and 4 assess potential differences between new brand name and new generic drugs. With prescriber fixed effects in column 4, we find that the estimated effect for generics ($Database_{ij} \times Generic_j$) remains positive and significant, while that on branded products is indistinguishable from zero. Specifically, physicians using the database are 2.08 percentage points more likely to prescribe a new generic within its initial year, but are no faster in the case of new branded drugs.[34] That database use may tilt prescribing toward faster generic adoption has potentially significant aggregate cost implications given the size of the market for cholesterol therapies and the chronic nature of the condition they treat.[35] Moreover, because generic drugs share identical clinical attributes with branded versions, these results strongly suggest database users may be responding to the increased salience of non-clinical information—in particular, price and insurance formulary data.

While the estimates in columns 1–4 rely on the full sample of US physicians, these results may be sensitive to underlying differences across database users and nonusers in the evolution of prescription outcomes. To better isolate the within-doctor impact of database use, columns 5 and 6 restrict the physician sample, including only those that both adopt and use the drug database during the sample period. Doctors that have yet to use the database by December 2010 are thus omitted. With this restriction, the estimates in column 6 suggest, as in column 4, that database users are indeed significantly faster to begin prescribing new generics, but show no significant effects on the adoption of new branded drugs.

Across all columns in Table 4, zip-code-drug fixed effects absorb variation across locations and over time in (i) access to other drug information sources (e.g., advertising), (ii) physicians' tendency to adopt new technology, (iii) patient characteristics affecting the price or match quality of drug j, and (iv) competition. Any component of factors (i)–(iv) that is stable over time is further captured by physician fixed effects in columns 2, 4, and 6. Comparing columns 3–6, however, it is evident that failing to include physician fixed effects and restrict the sample to eventual adopters tends to result in larger estimated associations between prescription outcomes and database adoption. Henceforth, we therefore present

[33] For comparison, we replicate Table 4 using $\ln T_{ij}$ as the outcome variable. Truncation in T_{ij} is addressed by limiting the duration of analysis to a window of 54 months following each new drug introduction and omitting Simcor; 54 months is the time span between the release of simvastatin and December 2010. We also consider $P\{T_{ij} \leq 6\}$ and $P\{T_{ij} \leq 24\}$ as alternatives. In each case, we find results that are qualitatively identical to those in Table 4.

[34] Notice that the zip-code-drug fixed effects ensure that this result is not explained by differences in local mandatory substitution laws. We nevertheless provide additional results regarding the effects of substitution laws on prescribing outcomes in Section IV.

[35] See Section VII.

Table 5—Prescription Diversity and Propensity, US Physicians, 2000–2010

Dependent variable	Number of unique drugs	Prescription HHI	$\mathbf{1}\{$(prescriptions of drug j by i at t) $> 0\}$
		Eventual users	
	(1)	(2)	(3)
$Database_{it}$	0.0350	−0.0027	
	0.0094	0.0010	
$Database_{it}$			
$\times New_{jt} \times Generic_j$			0.0158
			0.0025
$\times New_{jt} \times Branded_j$			−0.0051
			0.0015
$\times Old_{jt} \times Generic_j$			0.0240
			0.0029
$\times Old_{jt} \times Branded_j$			0.0013
			0.0009
$Prescription\ Volume_{it-1}$	0.0187	−0.0007	0.0013
	0.0002	0.0000	0.0000
Physician fixed effects	Yes	Yes	Yes
Zip-code-month fixed effects	Yes	Yes	Yes
Physician × t trends	Yes	Yes	Yes
Drug-month fixed effects	No	No	Yes
Observations	3,013,241	3,013,241	7,674,288
R^2	0.8941	0.7484	0.5458

Notes: This table provides least-squares estimates of (i) equation (6) in columns 1–2, and (ii) equation (7) in column 3, for cholesterol drug prescriptions by US physicians during January 2000 through December 2010 and the subset of physicians that eventually adopt and use the electronic reference to search for information about cholesterol drugs. Full-sample estimates appear in Table A.2 in the online Appendix. The dependent variable in column 1 captures the prescription diversity of physician i as the number of unique drugs j that are prescribed by i during month t. The dependent variable in column 2 is the prescription Herfindahl-Hirschman index for physician i in month t. The dependent variable in column 3 is an indicator for whether the doctor i prescribes drug j during month t. Database is the "Drug database and use indicator" variable described in Table 2, and takes a value of 1 for a physician user with database access in month t; it is otherwise zero. All regressions include the cholesterol drug prescription volume for physician i in month $t-1$, physician-specific time trends, and physician and zip-code-month fixed effects; column 3 also includes drug-month fixed effects. For computational ease, the estimates in column 3 rely only on observations in January 2000 and every subsequent June and December. Results in column 1 are robust to Poisson estimation; all columns are robust to including the first lag of the dependent variable. Results in column 3 are robust to logistic estimation, including doctor-drug fixed effects, and including zip-code-drug-month fixed effects. Standard errors clustered by zip code appear below each point estimate; results are robust to clustering errors by physician.

specifications that, like column 6, both include physician fixed effects and restrict the sample to eventual adopters.[36]

B. *Prescription Diversity*

To evaluate the impact of physicians' database use on the diversity of prescribing, Table 5 provides estimates of (6) for two outcome variables: the number of unique drugs M_{it} prescribed by physician i at t (column 1), and the associated Herfindahl-Hirschman index, HHI_{it} (column 2). The coefficient of interest is on

[36] All results described in this section are replicated in the online Appendix for the "All physicians" sample considered in columns 1–4 of Table 4.

$Database_{it}$ (Z_{it}), which takes a value of 1 if physician i has access to the drug database at t, and is otherwise zero.

In the model, prescribing diversity increases when a physician adopts the drug database if the prescription probability P_{ijt} rises more, on average, than it falls—that is, if adoption induces an increase in ϕ_{it} and if $\sum_{j=1}^{J_t} \partial P_{ijt}(\phi_{it})/\partial \phi_{it} > 0$. The results in Table 5 are strongly consistent with this. In column 1, the estimated coefficient on physician-i database access Z_{it} is indeed positive and highly significant, and indicates that database users prescribe, on average, 0.035 additional unique drug varieties each month relative to a nonuser.[37] In column 2, we find that HHI_{it} is also strongly responsive to database adoption; the estimate -0.003 is negative and highly significant, indicating that database users' prescribing is substantially less concentrated across drugs j. Given the assumptions of the model, these diversity results are consistent with the idea that database use is associated with increased prescribing *quality*, provided that the common component v_{jt} in (1) is small relative to ϵ_{njt}. Consistent with the model, we also find that prescription diversity increases significantly in a physician's monthly prescription volume.[38]

In both columns 1 and 2, physicians differ only in their respective drug reference adoption dates. Moreover, both specifications include doctor and zip-code-month fixed effects, as well as doctor-specific time trends. Nevertheless, the estimated coefficients are smaller than the corresponding full-sample estimates in online Appendix Table A.2, columns 1–2. This suggests the possible influence of *nonlinear* dynamic unobserved factors, correlated with Z_{it}, that influence prescription diversity. We discuss this concern in more detail in Section V below.

C. *Prescription Propensity*

We report estimates of (7) in column 3 of Table 5 for eventual database users; corresponding estimates for all physicians appear in online Appendix Table A.2. The outcome variable $\mathbf{1}\{X_{ijt} > 0\}$ is binary, indicating whether physician i writes at least one prescription for drug j at t, and the main coefficients of interest β_0 and β_1 correspond to $Database_{it}$ (Z_{it}) and its interaction with indicators for new generic and branded drugs, respectively.[39] We define j to be new while it is within $\tau = 24$ months of its initial market approval.

As described above, the model predicts that the coefficients β_0, β_1, β_2, and β_3 could be positive or negative depending on the distribution across drugs in the unobserved

[37] Because M_{it} is a count variable, we reestimate the coefficients in column 1 using a Poisson estimator and find qualitatively identical results.

[38] An interesting consideration is whether the model predicts, for any observable outcome variable, *no effect* of database use. One such "placebo" variable is the doctor-i prescription volume in month t. We reevaluate (6) using the number of prescriptions by doctor i in month t as the dependent variable and a sample identical to that in Table 5, columns 1 and 2. The estimated coefficient on database adoption Z_{it} is 0.0170 (standard error = 0.0301), which, in line with the model, is not statistically distinguishable from zero.

[39] The estimates in column 3 of Table 5 span January 2000 through December 2010. To manage the computational burden of the sample size, we include only observations in January 2000 and each subsequent June and December in estimating the coefficients. The results are not qualitatively or quantitatively sensitive to this sample size reduction.

quality terms v_{jt} and ϵ_{njt} about which physicians learn as ϕ increases. Note that P_{ijt} will tend to increase in ϕ_{it} for drugs with high v_{jt} values, but will generally fall if v_{jt} is relatively low. Whether database adoption increases or decreases, prescribing for a particular drug type is thus an empirical question. It is, however, one to which we can partially predict the answer, given the model and results in Table 4 and columns 1 and 2 of Table 5. First, Table 4 indicates database users are faster to begin prescribing new generics, but mildly slower to adopt new branded drugs; if the model is correct, this suggests, through (3), that $\beta_0 > 0$ and $\beta_1 < 0$. Second, prescription diversity increases with database adoption in columns 1 and 2 of Table 5; in (4), this implies that the sum across β_0, β_1, β_2, and β_3, weighted by the number of drugs per category, is positive.

In column 3 of Table 5, it is clear that physicians with database access are substantially more likely to prescribe generic products, regardless of vintage: $\beta_0 > 0$ and $\beta_2 > 0$. Specifically, these estimated coefficients are positive and highly significant, indicating the likelihood physician i prescribes a generic drug is 4.0 percentage points higher when she has access to the database. The coefficients on branded drugs are much smaller in magnitude; users are 0.13 percentage points more likely to prescribe an old, branded product ($\beta_3 > 0$), though the coefficient is not precisely estimated, and are 0.51 percentage points *less* likely to prescribe a new, branded product ($\beta_1 < 0$). Viewed through the lens of the model, the estimated coefficients $\hat{\beta}_0$ and $\hat{\beta}_1$ are thus exactly aligned with the results in Table 4. Moreover, the coefficients in column 3 are consistent with the diversity estimates in columns 1 and 2. Accounting for the distribution of drug types in the data, the weighted sum across the coefficients β in column 2 is always positive. The predictions of the model, in light of the estimates in Table 4 and columns 1 and 2 of Table 5, are thus strongly consistent with the results in column 3 of Table 5.

Importantly, these latter estimates provide additional suggestive evidence regarding the mechanism through which database adoption may influence prescribing. That the impact of database access is apparent not only for new, but also older products suggests that there could be important reasons for doctors to continually reference the drugs in question—possibly to learn about aspects of a drug that are either time-varying or patient-specific. This ongoing process of information acquisition could thus be an important factor in explaining why a significant fraction of a new drug's diffusion occurs beyond its first two years (Figure 2), a general feature of technology diffusion processes that has fascinated economists for decades (Manuelli and Seshadri 2014). Second, the estimated coefficients of interest are again consistent with responsiveness to economic information—price and formulary status—the inclusion of which is a distinctive feature of the drug database we consider. In particular, users tilt prescribing away from new, branded products—for which prices are generally high—and strongly toward generic products, for which prices are low. That the estimated increases in generic propensity (β_0, β_2) are so large in magnitude relative to those for branded products further suggests the response to price information is economically important. This is not only surprising, but also indicates database adoption could therefore have significant implications for aggregate prescription costs.

IV. Robustness and Alternative Specifications

The results in Tables 4 and 5 above indicate that physicians using the point-of-care reference begin prescribing new generics sooner than nonusers, prescribe a more diverse set of drugs than nonusers, and more often prescribe a generic drug. Because database adoption is not randomly assigned, however, a key question that remains is whether a causal interpretation of our results is supported by the data. Our sample restriction to eventual adopters and fixed-effects estimation approach helps rule out certain alternative explanations including cases in which "early-adopters" or physicians facing intense local advertising exposure begin both using the database and prescribing a new drug sooner than other physicians—for reasons unrelated to the actual impact of database use. A limitation of this approach arises when the timing of database adoption is either correlated with or involves selection on time-varying, idiosyncratic physician characteristics that are relevant to the prescribing outcomes we consider.

In the case of advertising, for example, if an individual physician's database adoption timing were to be correlated with the idiosyncratic change in her own exposure to drug detailing—that is, the component not accounted for by zip-code-month effects or doctor-specific time trends—our estimates of the impact of adoption could be confounded with the influence of drug advertising. Without granular data on detailing, it is not possible to know whether there are meaningful shocks to detailing efforts that differ substantially across physicians within a US zip-code-month for the drugs we consider.[40] It is important to note, however, that our results indicate database adopters are primarily quicker to prescribe new generic drugs, which are not advertised.[41]

Relatedly, it is possible in Table 5 that a physician's decision to adopt the database is partially determined by the underlying, idiosyncratic rate of increase in her prescribing diversity. Naturally, adopting the database today could be a more attractive option for a physician who anticipates prescribing a wider range of products in the future than for a physician in the opposite situation.[42] In this section, we describe additional results in order to clarify the degree to which a causal interpretation is warranted and to determine which alternative interpretations can be ruled out.

[40] Historically, detailing, or drug marketing efforts directed at individual physicians through sales representatives, has been the pharmaceutical industry's main promotional instrument (e.g., Manchanda and Honka 2013). As part of the detailing process, pharmaceutical sales representatives often provide doctors with payments related to drug promotion. IMS Health (IQVIA) does not curate doctor-specific data on detailing, though Verispan did in the past (Datta and Dave 2017); however, data on drug promotion payments by doctor, drug, and date are available from ProPublica for the period August 2013–December 2016. Although the latter data do not overlap with our sample period, precluding a direct consideration of detailing in our analysis, they are to our knowledge the best available information on doctor-drug specific detailing.

[41] The data studied by Larkin et al. (2017) show that detailing is almost exclusively done for branded drugs.

[42] While US physicians always face the same set of drugs approved for prescription, and in that sense do not differ in changes to the range of products available, their exposure to such changes may differ due to the potentially distinct characteristics of the specific patients they treat. To the extent that these distinctions are time-varying and correlated with database adoption, even across doctors practicing within the same zip-code-month, our fixed effects estimator could yield biased estimates.

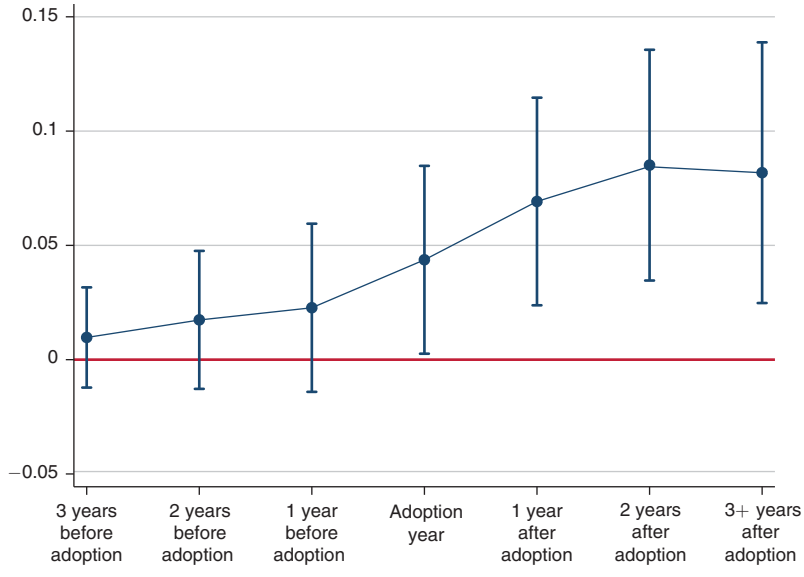

FIGURE 4. COEFFICIENTS BY YEAR FROM DATABASE ADOPTION—NUMBER OF UNIQUE DRUGS

Notes: This figure plots coefficients from a version of equation (6) that replaces the Database indicator with dummies corresponding to years before and after database adoption. Error bars show 95 percent confidence intervals. The full set of coefficients is shown in the online Appendix.

A. *Timing the Impact of Database Adoption*

For an important class of alternative explanations, physicians either adopt the reference database in response to pre-existing changes in prescribing, or are influenced by omitted, dynamic factors that simultaneously affect both prescribing and database adoption. A symptom that would likely appear where these influences are active is adopters who exhibit a trend toward the predicted outcomes, even before accessing the database.

While the physician-specific time trends included in (6) and (7) help to address this possibility, Figures 4, 5, and 6 examine the timing of the relationship between reference database adoption and changes in the (time-varying) prescription outcomes considered in Table 5; corresponding estimates appear in online Appendix Table A.3. As in Table 5 and Dranove et al. (2014), we focus this evaluation on the sample of doctors that eventually adopts and uses the database, and exploit variation across doctors in the year of adoption. Specifically, we replace our measure of database adoption Z_{it} in (6) with dummies for three years before adoption, two years before adoption, one year before adoption, the adoption year, one year after adoption, two years after adoption, and three or more years after adoption. The base period is four or more years before adoption.

The resulting estimates indicate a positive effect of database use on the number of unique drugs prescribed, M_{it}, in all years following the adoption year, with no statistically significant effects in years prior to the adoption year (Table 4).

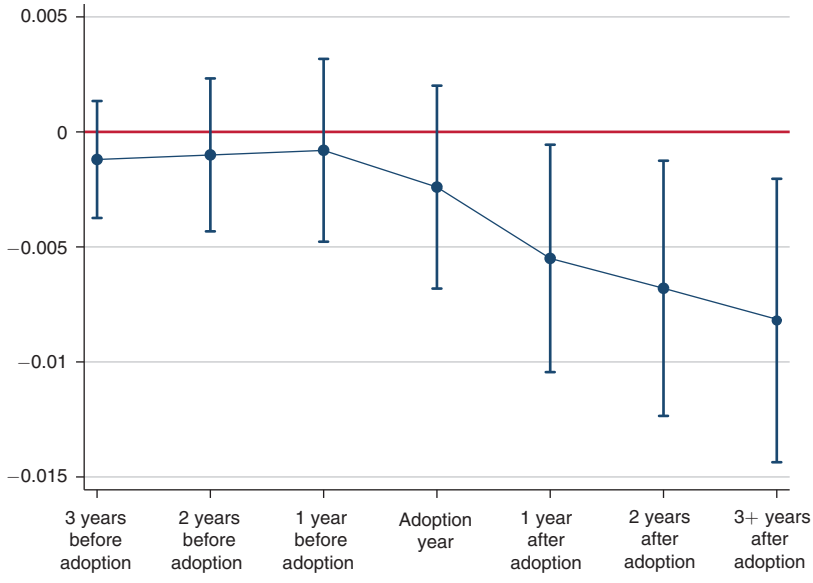

FIGURE 5. COEFFICIENTS BY YEAR FROM DATABASE ADOPTION—PRESCRIPTION HHI

Notes: This figure plots coefficients from a version of equation (6) that replaces the Database indicator with dummies corresponding to years before and after database adoption. Error bars show 95 percent confidence intervals. The full set of coefficients is shown in the online Appendix.

Similar patterns are shown in Figures 5 and 6 for prescription diversity measured as HHI_{it} and the generic prescription share, respectively. Taken together, Figures 4 through 6 indicate there were no clear trends in omitted variables that could have been driving the estimates in Table 5.

B. *Intensity of Use*

The model proposes a specific mechanism: physicians' prescription outcomes are influenced by database access because the information obtained through the database is important, yet otherwise unknown. If this proposed mechanism is correct, then for physicians using the database to search for cholesterol drugs with different intensities, those searching more intensely should have correspondingly larger prescription responses. There could certainly be other explanations linking search intensity with the prescribing outcomes we consider, but this evaluation is informative nevertheless: if intensity and response are *not* linked in the data, it would strongly suggest the proposed mechanism is invalid.

With these considerations in mind, we make use of a key feature of the data that allows us to observe not only a physician's database registration date, but also a proxy for the extent of her drug-month-specific database use, conditional on adoption. The variable is unfortunately not an exact lookup count, as the database company changed the way it maintained lookup data over time; given this, we aggregate the lookup proxy into a coarser measure reflecting a lower bound for each physician's overall intensity of database use. We then divide the sample of physicians into

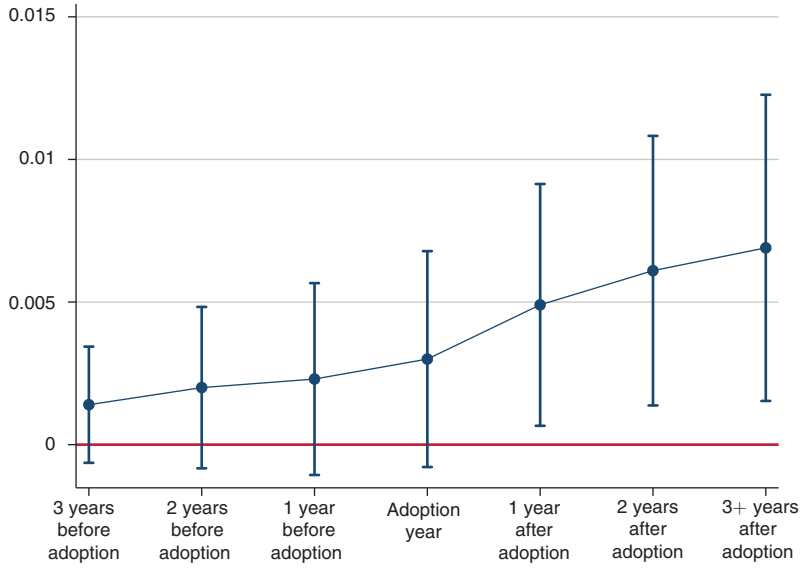

FIGURE 6. COEFFICIENTS BY YEAR FROM DATABASE ADOPTION—GENERIC PRESCRIPTION SHARE

Notes: This figure plots coefficients from a version of equation (6) that replaces the Database indicator with dummies corresponding to years before and after database adoption. Error bars show 95 percent confidence intervals. The full set of coefficients is shown in the online Appendix.

three groups based on this lower bound, and reassess the results in Tables 4 and 5 allowing different coefficients for each intensity group: low-intensity users are database adopters for whom the lower bound on cholesterol drug lookups is zero; the high-intensity group includes doctors whose total lookup proxy is in the top decile. The results appear in Tables A.4 and A.5 in the online Appendix.

The estimates in online Appendix Table A.4, column 1, reveal that the impact of database adoption is systematically and monotonically increasing in the intensity of usage. Low- and medium-intensity users are not more likely to adopt new generic drugs within a year of release, while high-intensity users are 2.36 percentage points more likely to do so; this latter coefficient is strongly significant, revealing meaningful heterogeneity across database adopters in its prescribing impact. Column 2 reveals a similar pattern of increasing impact in prescription diversity, although the pattern weakens for the prescription HHI in column 3.

Online Appendix Table A.5 suggests that the highest-intensity users are the most responsive to database access in switching away from branded drugs and toward generic products. The estimated coefficients $\beta_0 > 0$ and $\beta_2 > 0$ are approximately ten times larger for physicians searching the database intensely, relative to those with low-intensity database usage; as a result, an intense database user is 9.0 percentage points more likely to prescribe a generic in a given month than an adopter with low database usage. Moreover, the effects on branded varieties suggest additional differences: intense users are 1.9 percentage points less likely to prescribe branded varieties, while by contrast, adopters with low usage are actually 1 percentage point *more* likely to prescribe a branded drug per month.

These differential effects are consistent with the information mechanism we propose in the model, and suggest that the estimates in Table 5 are not simply capturing nonlinearly time-varying unobserved doctor characteristics that are on average correlated with database adoption; specifically, adopters in online Appendix Table A.5 with the lowest levels of database usage display only a minimal association between adoption timing and prescription outcomes. The fact that additional effort by an adopting physician is required to obtain a large prescription impact is consistent with the nature of the drug database, which as a searchable reference, can only impact prescribing to the extent that it is used.

C. *Endogenous Database Adoption*

Another approach to handling the endogeneity of database adoption is to find an instrument that generates quasi-random variation in a physician's database adoption decision, and to estimate the impact of information access relying on variation in this instrument. We found that the doctor-month specific share of other local physicians that have adopted the reference database is a robust predictor of adoption for a given doctor i at t. Since this share is plausibly uncorrelated with physician-specific unobservables that influence adoption decisions, we reassess the results in Tables 4 and 5 using this share as an instrument for database adoption. Estimates appear in Table A.6 in the online Appendix; a detailed discussion of the specification, validity, and mechanisms appears in Section A.3 of the online Appendix.[43] If our main results were driven by an omitted variable, instrumenting for database adoption should make its apparent effect on prescriptions disappear. But if anything the instrumental-variables estimates suggest the results in Tables 4 and 5 tend to understate the impact of database adoption. In online Appendix Table A.6, column 1, database users are 8.0 percentage points more likely to begin prescribing a new generic within its initial year (compared with 1.3 percentage points in column 6 of Table 4), relative to a nonuser, with no significant effects among new branded drugs. Columns 2 and 3 similarly suggest that the true impact of database adoption on diversity is an order of magnitude larger than the estimates in Table 5, while column 4 suggests Table 5, column 3 understates the impact of the database on prescribing by a factor of four for new generics (β_0) and a factor of ten for old generics (β_2). While these estimates suggest our results are not merely reflecting the endogeneity of database adoption, it is important to note that our leave-out mean instrument relies on variation in group composition that in many applications leads to small-sample bias from weak instruments, and that could confound interpretation in certain cases (Angrist 2014).

[43] First-stage estimates appear in Tables A.7 and A.8 in the online Appendix.

D. *Mandatory Substitution Laws*

To encourage cost savings, many US states impose regulations mandating generic substitution where available; in most cases, such laws have been in force since the 1970s (Grabowski and Vernon 1979) and were thus in effect during the sample period. This means that pharmacists dispense the generic version of a drug even if the physician prescribed the branded version. Since our data are collected from pharmacies and are based on prescriptions dispensed, a possible concern is that the patterns in our data reflect pharmacists' behavior rather than physicians' prescribing behavior—for example, if the implementation of mandatory substitution laws differs across states and over time in a manner correlated with physicians' adoption of the drug database.

On this point, it is important to note that all of our baseline results include zip-code-drug or zip-code-month fixed effects that absorb any impact of mandatory substitution laws on the prescription outcomes we consider—even where these laws may be correlated with physicians' database adoption choices. Nevertheless, to further check the robustness of our results to this potential concern, we reestimate the regressions reported in Table 4, replacing the zip-code-drug fixed effects with a set of drug fixed effects, and splitting the sample between doctors in states with versus without mandatory substitution laws. We also consider triple-interacted specifications in which the effect of database adoption is allowed to depend on whether the doctor is in a mandatory substitution state. The results, shown in online Appendix Table A.9, indicate that database adoption is associated with mildly *faster* adoption of new generics for physicians in mandatory substitution states.

While at first glance surprising, this result is consistent with the model given the estimates on prescription diversity shown in Table 5, column 1. Intuitively, mandatory substitution laws should reduce relative differences across doctors in generic prescribing, including differences resulting from database use. Even if doctors not using the database would be more inclined to prescribe a given branded drug j, the pharmacy always dispenses its generic equivalent k in a state with a mandatory substitution law. Therefore, when considering as an outcome variable the within-molecule generic prescription share $X_{ikt}/\left(X_{ijt}+X_{ikt}\right)$, mandatory substitution laws should eliminate systematic differences across doctors. On this basis, one would expect the database not to have a significant impact on $X_{ikt}/\left(X_{ijt}+X_{ikt}\right)$ in states with mandatory substitution laws.

On the other hand, with multiple molecules as in online Appendix Table A.9, and when the outcome variable is an indicator for whether doctor i prescribes drug j within its first year on the market, this reasoning leads to the opposite conclusion: mandatory substitution laws increase the apparent effect of database usage with respect to generic drug adoption. Even if the generic version k is available and every pharmacy dispenses j as k, this substitution occurs only if doctor i *actually attempts to prescribe* j. The results in Table 5 indicate database use increases prescription diversity, including along the extensive margin (number of unique drugs prescribed per month). Thus, the likelihood of a prescription of j occurring for doctor i within a year of the generic version k's release is *significantly higher* if doctor i is a database user; consequently, in a state with a mandatory substitution law, where this

prescription of *j* is dispensed as *k*, one should expect a *higher* estimated effect of database use. This is indeed consistent with the estimates in online Appendix Table A.9; in column 3, database users in mandatory substitution states are nearly twice as likely to adopt a new generic within its first market year.

E. *Pharmaceutical Innovation*

Physicians practicing in locations known for pharmaceutical innovation may have access to frontier knowledge regarding pharmaceutical development and pricing, limiting the potential for the reference database we observe to influence prescribing decisions. Within the conceptual framework outlined in Section III, proximity to the frontier could imply physicians have initially high ϕ_{it} parameters that are either minimally or not responsive to database use. If so, database use has little potential to affect prescribing. The zip-code-drug or zip-code-month fixed effects included in the baseline specifications account for the innovativeness of a physician's local environment. To assess whether location-specific differences in innovativeness impact the mechanism, we therefore replace these effects with either drug or month fixed effects, and use patent data from the NBER US Patent Citations Data File (Hall, Jaffe, and Trajtenberg 2001) to measure the number of pharmaceutical patents granted between 1975 and 1999 by zip code. We then reevaluate the adoption-lag specifications in Table 4 separately for two samples corresponding to the top and bottom 5 percent across zip codes based on the number of pharmaceutical patents granted. We also consider interacted specifications that account for differences in pharmaceutical patenting across locations.

The estimates appear in online Appendix Table A.10, and suggest that physicians plausibly located near the knowledge frontier—that is, physicians in zip codes among the top 5 percent by drug patenting—indeed respond to drug information differently than their more distant peers. Specifically, the estimates indicate that while use of the database in the least-innovative locations is associated with a larger impact on the likelihood of new drug adoption within one year (column 2) it has no significant impact in the most innovative locations (column 1). Considering the full sample, column 3 indicates that the database speeds generic adoption on average, but has especially pronounced effects among the least-innovative locations that are likely to be far from the information frontier. Column 4 confirms this result using a continuous measure of local patenting. Innovative areas adopt generics more quickly regardless of database adoption, but on average, a physician using the database in these locations is significantly less responsive to the information in terms of new generic drug adoption. The database nevertheless has an independent effect, speeding the adoption of new generic drugs regardless of patenting.

F. *Other Robustness Checks*

We evaluate the results involving the count variable M_{it} in Table 5, column 1 using a Poisson estimator. We estimate specifications involving the binary variables in Table 4 and column 3 of Table 5 using logistic fixed effects regressions. We reevaluate Table 5, column 3 including doctor-drug fixed effects and zip-code-month

drug-type fixed effects. To allow for persistence in prescription outcomes, we also control for the first lag of each outcome in Table 5.[44] Each of these robustness checks reveals qualitatively similar results.[45]

Finally, a physician's decision to prescribe a generic drug may be related to the insurance coverage of her patient population. We therefore evaluate split-sample estimates based on whether physicians receive a high or low share of Medicare and Medicaid patients, relative to the privately insured; separately, we repeat this split-sample analysis, distinguishing physicians based on whether a high or low share of their patients pay for prescriptions with cash. In both cases, we find negligible differences across groups.

V. Physician Heterogeneity

As a final point, we consider whether the data are broadly consistent with the idea that incomplete information contributes to disparities in prescribing behavior across physicians. If indeed these disparities partially reflect systematic informational differences, then physicians sharing access to a common source of drug information, like the reference database in our study, should tend to exhibit observable homogenization relative to other doctors. We consider this possibility using a simple approach. We first assign each physician to one of two groups based on her database registration status in December 2010. Then, within each group, we measure the extent of prescribing heterogeneity across physicians: specifically, we determine the vector of prescription shares for each prescriber *i* in December 2010, and then compute the Euclidean distance between this physician-*i* vector and the average vector of prescription shares among physicians in her group (database users or nonusers).

These within-group similarity measures are reported in online Appendix Table A.11, panel A. The prescription shares of database users are indeed more homogeneous than those of nonusers. The average Euclidean distance between the physician-*i* prescription vector and the group-specific average is 0.152 for users and 0.176 for nonusers, and the difference (−0.0236) is highly statistically significant. Importantly, note that database users prescribe a significantly more diverse set of products than nonusers, as shown above in Table 5; the relative homogeneity of database users' prescribing patterns thus does not imply a loss of variation in therapies generally. Rather, the result implies that physicians who are connected to the same information source resemble each other more closely in spite of the fact that they tend to prescribe a more diverse set of drugs.

Of course, the fact that database users' prescribing patterns are less heterogeneous could reflect selection rather than any causal effect of information access. Indeed, panel B of online Appendix Table A.11 shows that *eventual* database users' prescribing exhibited greater homogeneity than nonusers' even in January 2000, before anyone was using the database. But the changes by group between 2000 and 2010, summarized in panel C, indicate that while (i) both groups (users and

[44] Structural persistence could arise in the presence of persistent patient-specific match quality ϵ_{njt} terms in the model, given the chronic nature of the relevant medical condition.

[45] Detailed results available on request.

nonusers) exhibit homogenization over time, with the average Euclidean distance declining by 0.052 for users and by 0.040 for nonusers, (ii) the difference in differences is also highly significant—i.e., significantly more within-group homogenization is observed among database users than among nonusers, even when controlling for physician fixed effects.[46]

Our data cannot definitively say whether the faster convergence for database users was directly caused by the database; unlike the regression analyses reported in Section IV, the results described in online Appendix Table A.11 are based on across-group comparisons rather than within-doctor comparisons, so the stronger trend toward homogenization among adopters could reflect other characteristics that are correlated with the decision to adopt. Nevertheless, we view these results as suggestive of the idea that database use could reduce disparities in care—an idea that merits exploration in future research.

VI. Conclusion

This paper has empirically examined how physicians' prescribing decisions are affected by access to a drug reference database at the point of care. Using a novel dataset that includes prescription choices and drug reference use for over 125,000 individual US physicians, we find that after adopting the database, users increase the likelihood of prescribing a generic drug, are faster to begin prescribing a newly-released generic, and yet also significantly increase the diversity of products prescribed each month. These results are consistent with the predictions of a simple, incomplete-information model of prescription choice, and are robust across specifications that control for physician and location-month unobserved prescribing determinants and treat the timing of physicians' database adoption as endogenous.

While the magnitude of database users' estimated shift toward generic drugs is modest at the prescriber level, the implied aggregate impact on drug spending is economically significant. US pharmacies filled approximately 170 million cholesterol drug prescriptions in 2010, for example, a year during which roughly 45 percent of sample physicians were users of the drug reference database, and during which the average price difference between branded and generic cholesterol-drug prescriptions was around \$94.[47] If 45 percent of these prescriptions correspond to database users, and if users' generic shares increase by even half a percentage point—approximately the magnitude of the measured effect of database use in our data—the implied annual cost savings of database usage would exceed \$35 million for cholesterol drugs alone.[48] If effects of the same magnitude apply to all drug classes, the implied savings would be on the order of \$1 billion annually.

[46] For clarity, panel C reports coefficients from a least-squares regression of the Euclidean distance to the mean D_{it} for doctor i at $t = \{\text{January 2000, December 2010}\}$ on an indicator I_{2010} for December 2010, its interaction $Z_{2010} \times I_{2010}$ with an indicator $Z_{i,2010}$ for physician-i database access in December 2010, and physician fixed effects.

[47] Assuming each prescription is for a standard 30-day supply, this estimated price difference based on Marketscan data for December 2009 is conservative; CMS data indicate the price gap corresponding to Medicare and Medicaid patients is substantially wider.

[48] To determine the relationship between physicians' generic prescription share and database use, we estimate a version of equation (6) that replaces M_{it} with the share of physician-i prescriptions in month t that are accounted for by generics. The estimated coefficient on Z_{it} is 0.0061 (standard error 0.00021).

More importantly, our study speaks to policy debates regarding the efficiency of US healthcare provision, particularly those concerning unwarranted disparities in the observed cost and quality of medical care (Wennberg and Cooper 1996) including that involving prescription drugs (Munson et al. 2013). Our results provide new, systematic evidence that information differences contribute significantly to treatment variation across US physicians, and suggest that connecting physicians to common, high-quality information sources has the potential to meaningfully increase the efficiency of health care delivery.

REFERENCES

Agha, Leila. 2014. "The Effects of Health Information Technology on the Costs and Quality of Medical Care." *Journal of Health Economics* 34: 19–30.

Agha, Leila, and David Molitor. 2018. "The Local Influence of Pioneer Investigators in Technology Adoption: Evidence from New Cancer Drugs." *Review of Economics and Statistics* 100 (1): 29–44.

Angrist, Joshua D. 2014. "The Perils of Peer Effects." *Labour Economics* 30: 98–108.

Arrow, Kenneth J. 1969. "Classificatory Notes on the Production and Transmission of Technological Knowledge." *American Economic Review* 59 (2): 29–35.

Athey, Susan, and Scott Stern. 2002. "The Impact of Information Technology on Emergency Health Care Outcomes." *RAND Journal of Economics* 33 (3): 399–432.

Attewell, Paul. 1992. "Technology Diffusion and Organizational Learning: The Case of Business Computing." *Organizational Science* 3 (1): 1–19.

Berndt, Ernst R., Robert S. Gibbons, Anton Kolotilin, and Anna Levine Taub. 2015. "The Heterogeneity of Concentrated Prescribing Behavior: Theory and Evidence from Antipsychotics." *Journal of Health Economics* 40: 26–39.

Black, Sandra E., and Lisa M. Lynch. 2001. "How to Compete: The Impact of Workplace Practices and Information Technology on Productivity." *Review of Economics and Statistics* 83 (3): 434–45.

Bloom, Nicholas, Luis Garicano, Raffaella Sadun, and John Van Reenen. 2009. "The Distinct Effects of Information Technology and Communication Technology on Firm Organization." NBER Working Paper 14975.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2012. "Americans Do IT Better: US Multinationals and the Productivity Miracle." *American Economic Review* 102 (1): 167–201.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *Quarterly Journal of Economics* 117 (1): 339–76.

Bresnahan, Timothy, and Shane Greenstein. 1996. "Technical Progress and Co-invention in Computing and in the Uses of Computers." *Brookings Papers on Economic Activity Microeconomics* 26: 1–8.

Bronnenberg, Bart J., Jean-Pierre Dubé, Matthew Gentzkow, and Jesse M. Shapiro. 2015. "Do Pharmacists Buy Bayer? Informed Shoppers and the Brand Premium." *Quarterly Journal of Economics* 130 (4): 1669–1726.

Brooks, J. Renée, Elizabeth A. Cook, Cole G. Chapman, Puttarin Kulchaitanaroaj, Elizabeth A. Chrischilles, S. Welch, and Jennifer Robinson. 2014. "Geographic Variation in Statin Use for Complex Acute Myocardial Infarction Patients: Evidence of Effective Care?" *Medical Care* 52 (S3): S37–44.

Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad. 2015. "What Does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics." NBER Working Paper 21632.

Brynjolfsson, Erik, and Lorin M. Hitt. 2003. "Computing Productivity: Firm-Level Evidence." *Review of Economics and Statistics* 85 (4): 793–808.

Cannon, Christopher P., Michael A. Blazing, Robert P. Giugliano, Amy McCagg, Jennifer A. White, Pierre Theroux, Harald Dairus, et al. 2015. "Ezetimibe Added to Statin Therapy after Acute Coronary Syndromes." *New England Journal of Medicine* 372: 2387–97.

Cannon, Christopher P., Eugene Braunwald, Carolyn H. McCabe, Daniel J. Rader, Jean L. Rouleau, Rene Belder, Steven V. Joyal, et al. 2004. "Intensive Versus Moderate Lipid Lowering with Statins after Acute Coronary Syndromes." *New England Journal of Medicine* 350: 1495–504.

Caselli, Francesco, and Wilbur John Coleman II. 2001. "Cross-Country Technology Diffusion: The Case of Computers." *American Economic Review* 91 (2): 328–35.

Caselli, Francesco, and Wilbur John Coleman II. 2006. "The World Technology Frontier." *American Economic Review* 96 (3): 499–522.

Castelli, William P., Robert J. Garrison, Peter W.F. Wilson, Robert D. Abbott, Sona Kalousdian, and William B. Kannel. 1986. "Incidence of Coronary Heart Disease and Lipoprotein Cholesterol Levels—The Framingham Study." *Journal of the American Medical Association* 256 (20): 2835–38.

Chernew, Michael E., Richard A. Hirth, and David M. Cutler. 2009. "Increased Spending On Health Care: Long-Term Implications for the Nation." *Health Affairs* 28 (5): 1253–55.

Coleman, James, Elihu Katz, and Herbert Menzel. 1957. "The Diffusion of an Innovation among Physicians." *Sociometry* 20 (4): 253–70.

Coleman, James S., Elihu Katz, and Herbert Menzel. 1966. *Medical Innovation—A Diffusion Study*. Indianapolis: Bobbs-Merrill Co.

Comin, D., and B. Hobijn. 2004. "Cross-Country Technology Adoption: Making the Theories Face the Facts." *Journal of Monetary Economics* 51 (1): 39–83.

Cooper, Zack, Stuart Craig, Martin Gaynor, and John Van Reenen. 2015. "The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured." Yale Working Paper ISPS15-027.

Crawford, Gregory S., and Matthew Shum. 2005. "Uncertainty and Learning in Pharmaceutical Demand." *Econometrica* 73 (4): 1137–73.

Cutler, David, Jonathan Skinner, Ariel Dora Stern, and David Wennberg. 2015. "Physician Beliefs and Patient Preferences: A New Look at Regional Variation in Health Care Spending." NBER Working Paper 19320.

Datta, Anusua, and Dhaval M. Dave. 2017. "Effects of Physician-Directed Pharmaceutical Promotion on Prescription Behaviors: Longitudinal Evidence." *Health Economics* 26 (4): 450–68.

David, Paul A. 1966. "Chapter 1—The Mechanization of Reaping in the Ante-Bellum Midwest." In *Industrialization in Two Systems: Essays in Honor of Alexander Gerschenkron by a Group of His Students*, edited by Henry Rosovsky. New York: John Wiley and Sons.

Dickstein, Michael J. 2018. "Efficient Provision of Experience Goods: Evidence from Antidepressant Choice." New York University Working Paper 18-17.

Dranove, David, Chris Forman, Avi Goldfarb, and Shane Greenstein. 2014. "The Trillion Dollar Conundrum: Complementarities and Health Information Technology." *American Economic Journal: Economic Policy* 6 (4): 239–70.

Duggan, Mark, and Fiona Scott Morton. 2010. "The Effect of Medicare Part D on Pharmaceutical Prices and Utilization." *American Economic Review* 100 (1): 590–607.

Escarce, Jose. 1996. "Externalities in Hospitals and Physician Adoption of A New Surgical Technology: An Exploratory Analysis." *Journal of Health Economics* 15 (6): 715–34.

Forman, Chris, Avi Goldfarb, and Shane Greenstein. 2005. "How Did Location Affect the Adoption of the Commercial Internet? Global Village vs. Urban Leadership." *Journal of Urban Economics* 58 (3): 389–420.

Fuchs, Victor R. 1974. *Who Shall Live? Health, Economics, and Social Choice*. New York: Basic Books.

Gawande, Atul. 2009. "The Cost Conundrum." *New Yorker*, June 1. https://www.newyorker.com/magazine/2009/06/01/the-cost-conundrum.

Grabowski, Henry G., and John M. Vernon. 1979. "Substitution Laws and Innovation in the Pharmaceutical Industry." *Law and Contemporary Problems* 43: 43–66.

Griliches, Zvi. 1957. "Hybrid Corn: An Exploration in the Economics of Technological Change." *Econometrica* 25 (4): 501–22.

Hall, Bronwyn H., Adam B. Jaffe, and Manuel Tratjenberg. 2001. "The NBER Patent Citations Data File: Lessons, Insights and Methodological Tools." NBER Working Paper 8498.

Hubbard, Thomas N. 2003. "Information, Decisions, and Productivity: On-Board Computers and Capacity Utilization in Trucking." *American Economic Review* 93 (4): 1328–53.

Kastelein, John J.P., Fatima Akdim, Erik S.G. Stroes, Aeilko H. Zwinderman, Michiel L. Bots, Anton F.H. Stalenhoef, Frank L.J. Visseren, et al. 2008. "Simvastatin With or Without Ezetimibe in Familial Hypercholesterolemia." *New England Journal of Medicine* 358: 1431–43.

Klenow, Peter J., and Andrés Rodríguez-Clare. 1997. "Economic Growth: A Review Essay." *Journal of Monetary Economics* 40 (3): 597–617.

Law, M.R., N.J. Wald, and A.R. Rudnicka. 2003. "Quantifying Effect of Statins on Low Density Lipoprotein Cholesterol, Ischaemic Heart Disease, and Stroke: Systematic Review and Meta-analysis." *British Medical Journal* 326 (7404): 1423–27.

Larkin, Ian, Desmond Ang, Jonathan Steinhart, Matthew Chao, Mark Patterson, Sunita Sah, Tina Wu, et al. 2017. "Association between Academic Medical Center Pharmaceutical Detailing Policies and Physician Prescribing." *Journal of the American Medical Association* 317 (17): 1785–95.

Ledford, Heidi. 2013. "Cholesterol Limits Lose Their Lustre." *Nature* 494: 410–11.

Manchanda, Puneet, and Elisabeth Honka. 2013. "The Effects and Role of Direct-to-Physician Marketing in the Pharmaceutical Industry: An Integrative Review." *Yale Journal of Health Policy, Law, and Ethics* 5 (2): 785–822.

Manuelli, Rodolfo E., and Ananth Seshadri. 2014. "Frictionless Technology Diffusion: The Case of Tractors." *American Economic Review* 104 (4): 1368–91.

Mozaffarian, Dariush, Emelia J. Benjamin, Alan S. Go, Donna K. Arnett, Michael J. Blaha, Mary Cushman, Sara de Ferranti, et al. 2015. "Heart Disease and Stroke Statistics—2015 Update." *Circulation* 131 (4): e29–322.

Munson, Jeffrey C., Nancy E. Morden, David C. Goodman, Luca F. Valle, and John E. Wennberg. 2013. *The Dartmouth Atlas of Medicare Prescription Drug Use*. Chicago: American Hospital Publishing, Inc.

Parente, Stephen L., and Edward C. Prescott. 1994. "Barriers to Technology Adoption and Development." *Journal of Political Economy* 102 (2): 298–321.

Phelps, Charles E. 2000. "Chapter 5—Information Diffusion and Best Practice Adoption." In *Handbook of Health Economics*, Vol. 1A, edited by Anthony J. Culyer and Joseph P. Newhouse, 223–64. Amsterdam: North-Holland.

Rossebø, Anne B., Terje R. Pedersen, Kurt Boman, Philippe Brudi, John B. Chambers, Kenneth Egstrup, Eva Gerdts, et al. 2008. "Intensive Lipid Lowering with Simvastatin and Ezetimibe in Aortic Stenosis." *New England Journal of Medicine* 359: 1343–56.

Skinner, Jonathan. 2012. "Causes and Consequences of Regional Variations in Health Care." *Handbook of Health Economics*, Vol. 2, edited by Mark V. Pauly, Thomas G. Mcguire, and Pedro P. Barros, 45–93. Amsterdam: North-Holland.

Skinner, Jonathan, and Douglas Staiger. 2007. "Technology Adoption from Hybrid Corn to Beta-Blockers." In *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, edited by Ernst R. Berndt and Charles R. Hulten, 545–70. Chicago: NBER, University of Chicago Press.

Solow, Robert M. 1956. "A Contribution to the Theory of Economic Growth." *Quarterly Journal of Economics* 70 (1): 65–94.

Solow, Robert M. 1987. "We'd Better Watch Out." *New York Times Book Review*, July 12, 36.

Ventola, C. Lee. 2014. "Mobile Devices and Apps for Health Care Professionals: Uses and Benefits." *Pharmacy and Therapeutics* 39 (5): 356–64.

Wennberg, David E., and John E. Wennberg. 2003. "Addressing Variations: Is There Hope for the Future?" *Health Affairs* 22 (S): W3-614-617.

Wennberg, John E., and Megan M. Cooper, eds. 1996. *The Dartmouth Atlas of Health Care*. Chicago: American Hospital Publishing, Inc.