

THE LOCAL INFLUENCE OF PIONEER INVESTIGATORS ON TECHNOLOGY ADOPTION: EVIDENCE FROM NEW CANCER DRUGS

Leila Agha and David Molitor*

Abstract—Local opinion leaders may play a key role in easing information frictions associated with technology adoption. This paper analyzes the influence of physician investigators who lead clinical trials for new cancer drugs. By comparing diffusion patterns across 21 new cancer drugs, we separate correlated regional demand for new technology from information spillovers. Patients in the lead investigator's region are initially 36% more likely to receive the new drug, but utilization converges within four years. We also find that superstar physician authors, measured by trial role or citation history, have broader influence than less prominent authors.

I. Introduction

ACROSS many industries, technology adoption exhibits spatial clustering (Comin, Dmitriev, & Rossi-Hansberg, 2012). In medicine, technology diffusion shows similar patterns, with notable clustering within interpersonal networks (Coleman, Katz, & Menzel, 1957), hospitals (Escarce, 1996), and geographic regions (Baicker & Chandra, 2010). The potential for geographic proximity to facilitate the spread of knowledge and innovations across individuals or firms has been recognized at least since Marshall (1890). It has long been argued that local opinion leaders play a key role in this diffusion process (Rogers, 1962), yet direct empirical evidence has been limited. If local opinion leaders have significant sway, the value of well-informed, prominent physicians could extend beyond their contributions to the scientific literature or the treatment of their own patients; influential doctors may shape the practice of medicine across their region.

The limited evidence on the role of opinion leaders is in part due to the challenges of empirical identification. Clustered technology diffusion within a network, organization, or region could be driven by common local demand or capacity for technology adoption rather than information spillovers. Furthermore, it is difficult to isolate the role of opinion leaders as a specific pathway for knowledge spillovers. This study takes on both of these challenges, analyzing whether new cancer drugs approved by the Food and Drug Administration (FDA) are adopted more rapidly in the geographic regions containing study authors of the pivotal clinical study

used in the FDA review process. First, by comparing diffusion patterns across many new drugs with different locations of study authors, we are able to separate local demand or taste for technology from the role of knowledge spillovers. Second, by comparing the influence of more and less prominent physicians as measured by academic citations and clinical trial role, we demonstrate that superstar authors have a substantially broader reach than their less prominent peers.

Understanding the determinants of technology adoption is of particular interest in the health care context, where new technologies are a key factor underlying both rising costs of care and improved health outcomes in the United States (Newhouse, 1992; U.S. CBO, 2008; Smith, Newhouse, & Freeland, 2009). Against the backdrop of this aggregate growth is substantial heterogeneity across regions in the extent and speed of new technology adoption (Fisher et al., 2003a, 2003b; Skinner & Staiger, 2005) and large geographic disparities in access to new cancer treatments (Nattinger et al., 1992; Farrow, Hunt, & Samet, 1992; Fairfield et al., 2010; Bristow et al., 2014).

While previous research has documented an extensive role of social learning in determining technology adoption in developing countries (Conley & Udry, 2010; Adhvaryu, 2014), we may expect little role for local information frictions in this setting where adoption decisions are made by expert physicians with high human capital and ready access to scientific information. On the other hand, it is a setting with substantial uncertainty about the efficacy and appropriate applications of newly introduced drugs. As a result, the clinical trial authors' detailed knowledge of drug mechanisms, patient responses, and side effects may put them and their peers at an informational advantage in the early stages of a drug's diffusion.

Existing empirical studies on the role of geographic spillovers have primarily focused on the creative process of new ideas and technologies, such as the tendency for inventors to cite patents developed in their geographic region (Jaffe, Trajtenberg, & Henderson, 1993) or for academic citation patterns to follow migrant scientists (Azoulay, Zivin, & Sampat, 2012).¹ We expand on this work to investigate the geographic connection between research activity and the subsequent adoption of resulting technologies. Further, there is relatively little evidence outside of development economics (Banerjee et al., 2013) on how the prominence or connectedness of an opinion leader affects his or her influence. Our study sheds light on whether opinion leaders continue to matter in a context where technology adoption decisions are made by highly specialized experts.

Received for publication August 28, 2015. Revision accepted for publication January 30, 2017. Editor: Amitabh Chandra.

* Agha: Dartmouth College and NBER; Molitor: University of Illinois and NBER.

We thank Joshua Aronson, Brittany Bychkovsky, Amitabh Chandra, Joseph Doyle, Amy Finkelstein, Jonathan Gruber, Nolan Miller, Jim Rebitzer, and Laila Saied, as well as seminar participants at Boston University, Georgia State University, MIT, NBER, Northwestern Kellogg, and RAND for helpful comments and suggestions. We are also grateful to Jean Roth and Mohan Ramanujan for assistance obtaining and managing the data. This research was supported by the National Institute on Aging, grant T32-AG000186.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00670.

¹ Audretsch and Feldman (2004) provide an extensive review.

Our analysis is based on a novel data collection effort that identified the study authors of the pivotal clinical trials for 21 new cancer drugs and matched the locations of these authors to adoption patterns of the drugs using Medicare claims records from over 1.4 million patient cancer care episodes from 1998 to 2008. For scientific publications of pivotal clinical drug trials, the principal investigator is typically credited with the first author position. The last author is often a research scientist affiliated with the sponsoring drug company and is not a practicing physician for most of the drugs in our sample. We restrict attention to study authors who are also practicing clinicians.

The key finding from our baseline analysis is that patients treated in the hospital market where the first author is located are 36% more likely to receive treatment with the new drug within the first two years following a drug's FDA approval. We show that alternative definitions of superstar authors based not on authorship position but on citation counts to previous publications yield similar results.

Increased drug use in the first author's region is driven by higher rates of adoption by physicians both within and outside the author's practice group, suggesting the first author's influence extends beyond the boundaries of his or her organization. Other physician study authors boost utilization only within their own physician group; non-first authors have smaller effects on regional patterns of care.

While initial proximity effects are large, the effects fade over time so that there is no discernible effect four years after a drug's approval. Despite this eventual convergence, initial differences in new drug use have significant implications for access to care, which we explore in section IV of this paper. This pattern of convergence also sheds light on the underlying driver of regional disparities in our context. The Roy model of medical treatment choice by Chandra and Staiger (2007) demonstrates that in the presence of productivity spillovers, greater physician experience with a treatment may lead to steady-state specialization in that treatment relative to alternatives, but our findings suggest that experience-related productivity spillovers are not a key factor in explaining regional disparities in cancer treatments.

We also examine the impact that author location may have on where patients seek care. We document that appropriate patients appear more likely to travel into a study author's region after a new drug is approved, suggesting that patients may benefit from access to broad provider networks. An instrumental variables strategy based on whether patients reside in an author's hospital market reveals that differential patient sorting accounts for about one-third of our main finding.

Finally, we probe the welfare implications of our findings by studying the survival improvements associated with the adoption of new cancer drugs. By comparing regions with fast and slow drug diffusion tendencies before and after the introduction of a new drug, we estimate that fast-diffusing regions show evidence of higher returns to new drug use. The survival improvements are so large that they appear unlikely

to be driven solely by the greater fraction of patients receiving the new treatment in the fast-adopting regions. Rather, they point to larger average treatment effects that could result from better physician selection of patients for treatment or improved dosing. This evidence further supports the idea that the local information environment may be a key determinant of both adoption and returns to new drug use and suggests that policies that boost utilization without changing the quality of local information may fail to realize the full potential benefits of the new technology.

The organization of the paper is as follows. Section II describes the empirical context and key data elements. Section III lays out the primary empirical strategies and results. Section IV investigates the role and extent of patient travel and selective sorting. Section V describes evidence on the survival benefits of new cancer drug adoption, and section VI concludes.

II. Setting and Data

In the United States, prescription drugs are regulated by the U.S. Food and Drug Administration (FDA), which between 2004 and 2013 granted approval for 26 new drugs per year on average (U.S. FDA, 2014). In order to receive approval, new drugs undergo an extensive review process, in coordination with both the drug's sponsor (the manufacturer) and the FDA.² This process begins with the submission of an Investigational New Drug application, which includes a proposal for testing the drug on human subjects through clinical trials. FDA regulations place primary responsibility on the sponsor to select clinical investigators and research sites, and on a qualified institutional review board to review and approve each clinical investigator's qualifications before participation in the investigation (21 CFR § 312.53).

While each drug application may cite several studies from various stages of drug development, the applicant must prespecify a pivotal trial, which is typically a randomized controlled trial that provides the most comprehensive evidence to date on the efficacy of the drug. For establishing generalizable efficacy of treatment effects, FDA industry guidance states that drug sponsors may wish to invoke a multicenter trial design since the results arise from a broader patient population and multiple clinical settings (U.S. FDA, 1998). Drug sponsors may further minimize the risk of an unsuccessful trial by employing a design in which the investigator enrolls a small fraction of the total number of subjects, especially in cases where the investigator has a disclosable financial interest in the study (21 CFR § 54.5(c)). As we will show, the pivotal trials for all the drugs we study utilized a multicenter trial design.

Once clinical testing is complete, the drug sponsor submits the results as part of a New Drug Application. If approval is granted, an official prescription drug label is

²Detailed information on the development and approval process for new drugs can be found on the FDA's website at <http://www.fda.gov/Drugs/DevelopmentApprovalProcess>.

written describing the indications for which the drug may be legally marketed. The FDA publicly releases this label as part of a detailed approval package describing the basis for approval and the pivotal clinical trial that provided the primary support for the approval decision.

Many new cancer drugs are approved based on promising evidence for a narrowly defined indication. For example, many clinical trials are conducted on patients whose cancers have relapsed after initial treatment, in which case the efficacy of the new drug as an initial treatment is not established upon drug approval. In addition, many cancer drugs come with side effects that range from temporary but severely uncomfortable (e.g., nausea, fever, pain) to serious or life threatening (e.g., kidney failure, lung damage, nerve damage, secondary cancers). A host of other drugs and additional monitoring may be required to mitigate these side effects, and physicians may develop expertise in this management over time.

In this study, we investigate the adoption and utilization of new cancer drugs in the years following initial approval by the FDA. Two key data elements are necessary for our analysis: the utilization of new cancer drugs across regions over time and the location of physician study authors who lead the pivotal clinical trials on which each drug's initial FDA approval was based.

A. Measuring Cancer Drug Use

We measure the diffusion of cancer drugs using Medicare Part B reimbursement claims over the eleven-year period 1998 to 2008. During the study period, 21 new cancer drug agents covered by Medicare Part B were approved by the FDA.³ The diffusion of these drugs forms the basis of our analysis.

While Medicare Parts A and B do not pay for most outpatient prescription drugs, an exception is made for drugs that are not typically self-administered, including cancer drugs administered intravenously or intramuscularly. These payments have comprised a rising proportion of Medicare spending in recent years. In 2004, Medicare Part B spent \$11 billion on drugs, a category dominated by cancer drug expenses; these costs rose 267% in the seven-year period since 1997, as compared to a 47% rise in total Medicare spending (Bach, 2009). Medicare Part B drug spending also comprises a significant share of total Medicare drug spending. As of 2010, spending on Part B drugs totaled \$19.5 billion, compared to the \$61.7 billion spent on Part D drugs, which are typically self-administered (U.S. Government Accountability Office, 2012).

We analyze drug use at the level of Hospital Referral Regions (HRR), as defined by the *Dartmouth Atlas of Health Care*, which partition U.S. postal code areas into 306 distinct

³ While more than 21 new cancer agents were introduced over this period, only drugs that are typically administered in a doctor's office are billable through Medicare Part B. Because we can track drug utilization only with Medicare Part B claims, this restricts our sample to this group of 21 agents.

FIGURE 1.—LOCATIONS OF DRUG PIVOTAL STUDY AUTHORS



Circles mark the Hospital Referral Regions (HRR) that contain a first author for the pivotal clinical trial of a cancer drug in our sample. Some regions are the site of multiple first authors: Houston (four first authors); Chicago (three first authors); Durham, North Carolina (two first authors); and New York, New York (two first authors). Triangles mark HRRs that contain other authors.

regions (see figure 1). Regions are defined by where the majority of the population in each postal code are referred for tertiary health care services and are commonly used as the unit of analysis for cancer care (see, e.g., Fisher et al., 2003a, 2003b; Onega et al., 2008).

To track the adoption and use of new cancer drugs, we analyze a 100% sample of Medicare outpatient claims, as well as a 20% sample of Medicare physician claims. For each drug in our sample, we study a cohort of patients diagnosed and treated for the targeted cancer type (e.g., colon, lung) for up to four years after the drug's initial approval. The unit of observation is the patient-year episode of cancer care; we analyze all claims associated with a cancer diagnosis for that patient within a calendar year. Our data comprise 1.4 million cancer care episodes within the first four years following drug introduction, of which 659,000 occur within the first two years after drug introduction. These data allow us to track the utilization rate of each new cancer drug among indicated patients across HRRs and over time following drug approval.

For this analysis, we identified 21 new cancer drugs that were covered by Medicare Part B and FDA approved between 1998 and 2007.⁴ Of these drugs, 17 had clinical trials led by researchers in the United States and thus may be used to identify the impact of proximity to a first author on drug diffusion. The remaining four drugs are included in the sample to improve the precision of coefficients on other control variables.

Summary information for the 21 drugs in our study and their pivotal clinical trials is listed in table 1, sorted in order of drug FDA approval date (further details on these drugs are provided in appendix table A1). These drugs target a variety of cancer types, including common carcinomas such

⁴ We began with a list of 26 new cancer drugs that we obtained from Bach (2009). Of these drugs, 5 of them were billed fewer than ten times in our sample over the first two years after approval. Given that we were not able to observe any measurable diffusion for these agents, they were excluded from the analysis.

TABLE 1.—LIST OF STUDIED CANCER DRUGS

Trade Name (1)	FDA Approval Year (2)	Target Disease (3)	Pivotal Clinical Trial				
			Trial Sites (4)	Patients Enrolled (5)	Publication Year (6)	Number of Trial Authors (7)	First Author City (8)
Xeloda	1998	Breast cancer	25	163	1999	10	Dallas
Herceptin	1998	Breast cancer	54	222	1999	11	Chicago
Valstar	1998	Bladder cancer	41	90	2000	6	Chicago
Ontak	1999	Cutaneous T-cell lymphoma	25	71	2001	26	Durham
Temodar	1999	Brain cancer	21	225	2000	22	Houston
Ellence	1999	Breast cancer	39	710	1998	18	Canada
Mylotarg	2000	Acute myeloid leukemia	32	104	2001	17	Seattle, WA
Trisenox	2000	Acute myeloid leukemia	9	40	2001	15	New York
Campath	2001	Chronic lymphocytic leukemia	21	93	2002	11	Houston
Zometa	2001	Hypercalcemia of malignancy	87	287	2001	11	Canada
Zevalin ^a	2002	Non-Hodgkin's lymphoma	45	200	2002	13	Rochester, MN
Faslodex	2002	Breast cancer	83	400	2002	14	Houston
Eloxatin	2002	Colon cancer	120	463	2003	8	Nashville, TN
Velcade ^b	2003	Multiple myeloma	14	202	2003	21	Boston
Bexxar	2003	Non-Hodgkin's lymphoma	3	40	2005	7	Stanford, CA
Alimta	2004	Lung cancer	88	456	2003	13	Chicago
Erbitux	2004	Colon cancer	56	329	2004	12	United Kingdom
Avastin	2004	Colon cancer	164	923	2004	15	Durham, NC
Dacogen	2006	Myelodysplastic syndromes	23	170	2006	16	Houston
Arranon	2006	Colon cancer	81	463	2007	12	Belgium
Torisel	2007	Kidney cancer	148	626	2007	19	Philadelphia

^aThere were two pivotal trials for Zevalin; the second trial had 11 authors, with the same first author also in Rochester, Minnesota.

^bThere were two pivotal trials for Velcade; the second trial had fifteen authors, with the first author in New York.

as breast, lung, and colon cancer, as well as hematologic and urologic cancers. Nearly all the pivotal trials for drugs in our study were large, multicenter trials. On average, each trial enrolled 299 patients across 56 trial sites. A majority of the pivotal clinical trials (13/21) were published in the *Journal of Clinical Oncology*, followed by the *New England Journal of Medicine* (4/21 publications).

Table A1, column 5 reports the number of indicated patient cancer care episodes observed in the two calendar years following FDA approval. There is substantial heterogeneity in target population size due to variation in disease prevalence, ranging from about 800 episodes of the relatively rare cutaneous T-cell lymphoma to over 84,900 episodes of lung cancer. An observation in our regressions is a patient cancer episode, effectively weighting our regressions by the size of the target patient population; however, the results from our baseline specifications are qualitatively unchanged when each drug is given equal weight.

B. Author Roles and Locations

In addition to the Medicare claims data, we also collected a new data set linking cancer drugs to the pivotal clinical trial that provided the primary support for FDA approval. By matching the pivotal trial information in the FDA application to the authors of the academic article reporting the trial's findings, we are able to identify the researchers who were primarily responsible for the trial.

There were an average of 14 authors per paper in our sample, ranging from 6 to 26 (table 1, column 7). We restricted our analysis to studying the influence of authors who are also practicing clinicians, excluding from analysis

the drug company–employed scientists who often coauthor clinical trials.⁵ We categorized authors as “first” and “other” according to the order listed on the clinical trial, and we recorded each author's location based on the postal code of the author's institution at the time of the article's publication. As indicated by table 1, there are many more trial sites than authors for most drugs in our sample; authorship typically signals intellectual involvement with the research process that goes beyond facilitating a clinical trial site.

Our analysis exploits the convenient fact that authorship order is a strong signal of author contribution and involvement. In contrast to other types of biomedical publications where the last author is often the principal investigator, the first author in large clinical trials is typically a senior physician who was leading the trial effort as the principal investigator (Baerlocher et al., 2007). The first author was a practicing clinician for all of the studies in our sample, in contrast to the last author position, which was held by a practicing clinician for only 7 out of 17 drugs. (The last author is frequently an employee of the sponsoring drug company.) The first author was also the single most highly cited clinical author on the trial in 8 out of 17 cases, but this is true for only 1 of the 7 last author clinicians.⁶

⁵In no cases was the first author of the study affiliated with the drug company. Often the drug company employees' contributions were credited with a middle or final author slot.

⁶As further evidence of the prominent role of first authors in this setting, a search of clinical trials listed on ClinicalTrials.gov found that for over 80% of the drugs in our sample, the first author is registered as an investigator for at least one trial involving that drug. By contrast, only 24% of last authors are listed as investigators for that drug. We do not use the ClinicalTrials.gov database to assign author roles because many of the clinical trials in our sample predate the formation of this registry, so it is not possible to match to

TABLE 2.—DRUG USE SUMMARY STATISTICS

Variables	First Author HRR (1)	Other Author HRR (2)	Author HRR for Different Drug (3)	HRR with No Authors (4)
Drug utilization rate	0.156	0.097	0.092	0.086
Fraction treated in author's group	0.534	0.357	0.000	0.000
Number of observations	6,985	29,322	250,330	372,831
Average number of patients per HRR per drug	388	236	254	75
Number of HRR-drug pairs	18	124	986	4958
Number of unique HRRs	11	54	54	252

Regions are defined by the 306 *Dartmouth Atlas Hospital Referral Regions* (HRRs). For each drug in the sample, regions are partitioned into four groups based on geographic proximity to authors of the pivotal trial, corresponding to the four columns in the table. Statistics are then reported for each column by aggregating over the set of drugs in the sample. Reported statistics reflect drug utilization over the first two years following initial introduction. Data on drug utilization come from Medicare claims, 1998–2008.

In addition to using authorship position to determine a physician's trial role and prominence, we also develop a second measure of each author's prominence within academic medicine using data on publication and citation histories from Web of Science. We rank authors based on citation counts accruing to publications in the relevant medical field. Specifically, we find all research articles matching each author's last name and initials over a ten-year period leading up to the year of FDA approval. Using keywords coded within Web of Science, we restrict to articles within the relevant field. For all authors in our sample, this includes oncology articles. Dermatology, neurology, hematology, urology, and nephrology are added to the definition of matched articles for drugs targeting those specific cancer types. These field restrictions provide a more targeted measure of prominence within the relevant medical field, as well as help disambiguate authors with common names. Next we count all citations that have accrued to those publications to the present. Finally, we define the top 10% (or 50%) authors as the top-cited author on that drug trial plus any other author whose citation count places him or her in the top 10% (50%) of all authors on the same drug's trial.

This citation-based measure of superstar status captures authors' academic prominence in their field, tagging the most highly cited authors on each trial. To disentangle superstar influence from differences across drugs or subspecialties in rates of drug take-up and prominence of investigators, we prefer this relative measure of superstar status, which allows us to compare the influence of authors within the set of researchers associated with each drug.

We categorize hospital referral regions (HRRs) based on their geographic proximity to clinical trial authors. For each drug, we define "first author HRR" and "other author HRR" to be the HRRs where the first author and other authors for that drug practiced, respectively. We create nonoverlapping definitions of author regions, so that a region cannot be coded as both a first and other author region, with first author designation taking precedence. Similarly, we code regions that contain superstar authors by the citation count metric.

The first authors on these trials practice at a wide set of academic medical centers. Author locations are pictured on

a map in figure 1, with first author locations marked by circles and other author locations by triangles. The most frequent first author locations within our sample are Houston (four first authors); Chicago (three first authors); Durham, North Carolina (two first authors); and New York (two first authors). There are 11 unique HRRs that contain a first author for at least one drug; 54 HRRs that contain a non-lead author for at least one drug but never contain a first author; and 252 remaining HRRs that never contain any author (see table 2). For our baseline results, we match patients to HRRs on the basis of where they received cancer care.

Within the HRRs that contain study authors, we further separate patients treated by doctors in the study authors' own practice groups from patients treated by other doctors in the region. To accomplish this, we group together all physicians who bill to the Medicare Carrier files using the same tax identification number as a clinical trial author. If a patient has at least one bill that year from a physician who is linked to a trial author's tax ID, we code the patient as treated within an author's practice group. Because academic oncologists typically work as part of large group practices, this allows us to separate the author's influence within his own organization from his or her influence on outside physicians.

C. Summary Statistics


The difference between a drug's utilization in the first author's region versus other regions can be seen in figure 2, which plots the fraction of indicated patients who receive the new drug for each year following FDA approval.⁷ The solid line plots raw drug utilization rates in each drug's respective first author region (HRR), and the dashed line plots the average drug utilization in all other regions.

Two of the main results from our empirical analysis that follows are suggested by the raw data plotted in figure 2. First, the figure shows that when a new cancer drug is introduced, indicated patients treated in the region containing the drug's first author are more likely to receive the drug than are patients in other regions. Second, this gap closes over time, so that drug utilization in the first author's HRR is no more intensive than in other regions after four years. This

the pivotal trial to registry information, even if subsequent trials involving the same drugs and investigators are part of the ClinicalTrials.gov database.

⁷ To make rates over time comparable, this graph is based on the balanced sample (67%) of drugs for which four years of usage rates are available in our data.

FIGURE 2.—NEW CANCER DRUG UTILIZATION RATES,
BY YEARS SINCE FDA APPROVAL



The figure plots raw usage for new drugs in the first author's HRR versus all other HRRs. The "First author HRR" and "Other HRRs" rates are generated by averaging over the drug-specific first author HRR and non-first author HRR rates calculated separately for each drug in each year. To make rates over time comparable, we restrict to the 67% of drugs in our sample for which four years of usage rates are available.

convergence occurs primarily because other regions increase their new drug utilization until it is similar to the first author region's use, not because use in the first author's region decreases substantially. If physicians are learning about the drug's value over time, it appears they are learning that the drug is a valuable addition to their practice. We do observe a slight decrease in average utilization rates in first author regions between years 3 and 4, from 18.0% to 17.2%, which could reflect new competing drugs entering the market.

More detailed summary statistics are reported in table 2. Over the first two years following drug approval, the average utilization rate of new drugs for indicated patient episodes ranges from 15.6% in regions where the first author practices to 8.6% in regions that never contain any investigators. Among the regions that do not contain any investigators for a given drug (columns 3–4), those regions that contain authors for other in-sample drugs (column 3) are the more intensive adopters, with 9.2% of patients receiving new drugs compared to 8.6% in regions that never contain any investigators. This suggests that authors tend to be located in regions that have a high degree of enthusiasm, expertise, or patient demand for new cancer drugs in general. Within the set of regions that contain an author for at least one in-sample drug, the first author HRR has 70% greater utilization rates on average (columns 1 and 3). Thus, despite the overall higher rates of new drug use in regions that contain an author, utilization is even greater when the lead researcher of the particular observed drug is in the area.

The second row of table 2 reports the fraction of patients treated in each region type who are ever treated by a physician in an author's practice group. Fifty-three percent of patients treated in the first author's region receive treatment from a doctor in the same practice group as a study author; only 36% of patients in a middle or last author's region

receive treatment from a doctor in an author's practice group. Given that the authors' practice groups do not have complete regional penetration for the targeted cancers, we can compare drug utilization within and outside the authors' own practice groups to test whether their influence extends beyond their own organization.

III. Empirical Evidence

A. Empirical Strategy

Our central idea is to exploit variation in the geographic location of lead study authors across multiple new cancer drugs to identify the impact of geographic proximity to these investigators on drug utilization. If the location of study investigators were randomly assigned across the country, we could simply compare drug utilization across locations and infer that any increased propensity to use the drug in an investigator's region was due to this proximity. However, study authors are not chosen randomly, and their locations are likely to be correlated with other regional factors (e.g., co-location with innovation-loving physicians) that influence the rate of new drug adoption.

The key methodological innovation in our analysis is to identify the effect of proximity to a drug's pivotal study authors by implementing an empirical design analogous to a difference-in-differences approach. Specifically, we compare drug utilization in study author and nonauthor regions, controlling for baseline differences in each region's propensity to use new cancer drugs, as well as controlling for time variation in drug utilization to capture the demand for each drug. Since we observe the diffusion of 21 newly introduced drugs, our strategy allows us to exploit each region's usage of other new drugs to establish its propensity to adopt a new drug when the region does not contain a study author. In addition, we use the time path of drug usage in nonauthor regions to establish how the drug usage evolved in the absence of author influence.

Our baseline regression specification takes the following form:

$$\begin{aligned}
 (drug)_{ijt} = & \beta_f \mathbf{1}(first\ author\ HRR)_{jd} \\
 & + \beta_o \mathbf{1}(other\ author\ HRR)_{jd} \\
 & + \{HRR \times disease\text{-}group\ FEs\}_{ijt} \\
 & + \{drug \times year\ FEs\}_{dt} + \delta X_{it} + \varepsilon_{ijt}. \quad (1)
 \end{aligned}$$

An observation is a patient-drug episode (patient i treated in provider region j , t years after drug d was approved), limited to episodes for which drug d is indicated based on patient diagnoses. The regression is estimated over patient-drug episodes that fall within two years following FDA approval of the drug.

The first two terms in the regression above are the key independent variables of interest: indicators for whether a study author of drug d 's pivotal clinical trial is located in region j . To reflect the possibility that proximity effects may

differ for these lead authors, we split our author proximity measures into separate indicators for whether a patient is treated in the first author HRR versus in a region containing any other author (but not the first author). The coefficients on these indicators describe how much more likely it is for a cancer patient to receive a new drug if treated in an HRR where an author of the drug's pivotal clinical trial is located.

The third term in this regression is a vector of fixed effects measuring each HRR's propensity to use new cancer drugs for each of three cancer disease types. Targeted diseases are grouped based on the cancer subtype: hematologic cancers (leukemias and lymphomas), urologic cancers (kidney and bladder cancer), and other carcinomas (brain, breast, colon, and lung cancer).⁸ This allows regions to differ in their enthusiasm and patient suitability for new cancer treatments within each disease group. The fourth term in the regression allows each drug to face an idiosyncratic yearly shock to utilization that is common across regions. Finally, we include patient characteristics X_i , which include indicators for patient sex, race, age (in five-year bins), and whether this is a new cancer treatment episode (i.e., patient had no cancer claims in the previous calendar year).

The primary threat to the validity of this approach stems from the possibility that study author regions are systematically more likely to use the new drug (for reasons not driven by author proximity) than their utilization of other new drugs for this cancer type and the national utilization of this particular new drug would predict. This threat could occur if, for example, clinical trials were located in areas with idiosyncratically high latent demand for that particular drug. As outlined in further detail below, we address this potential threat to validity in a number of other ways, including limiting the analysis to regions that ever contain a study author and studying the persistence of our estimated proximity effect. To preview our findings, the fact that the measured proximity effect converges within four years suggests that there are no permanent differences in patient appropriateness or latent demand for the new drug in first author regions.

In the first set of results discussed below and presented in section IIIB, we match patients to provider regions based on where patient care is delivered. Thus, any effect of author status on a region's propensity to prescribe a new cancer drug could be driven by two separate channels: (a) a prescribing effect in which providers in the author region have an increased propensity to treat a given population of patients

with the new drug; and (b) a sorting effect in which patients suitable for particular treatments (based on clinical appropriateness or patient demand) sort to providers who specialize in those treatments.⁹ For example, an increased number of suitable patients may travel into an author region for treatment, or suitable resident patients may be more likely to stay within the region for their care.

Our baseline specification in equation (1) measures the aggregate impact of first author status on drug utilization, but does not disentangle the prescribing and sorting mechanisms. Because these two channels have very different implications for policy, section IV applies an instrumental variables approach to isolate the change in utilization driven by prescribing—the increased propensity of first author regions to treat a given patient population.

B. Baseline Proximity Effects

Effects by geographic proximity. We begin by presenting evidence on whether geographic proximity to a new cancer drug's pivotal study authors affects a physician's propensity to prescribe that drug for indicated patients and, if so, how the proximity effect evolves over time.

With a Roy model of productivity spillovers, we may find geographic specialization in the use of medical treatments as described by Chandra and Staiger (2007). High-use areas develop expertise in the technology and have higher returns to its usage, and so they continue to use it more frequently in the steady state than low-use areas that do not develop a similar expertise. Under this model of productivity spillovers, we might expect to find long-run differences in the use of new cancer drugs across author HRRs and other regions. An alternative model such as Phelps (2000), where information asymmetries are the reason for delayed adoption among non-first author regions would predict convergence in practice patterns as information about the new treatments reaches each physician.

To measure the evolution of the author proximity effect, we estimate a modified version of specification (1) in which the author HRR indicators are interacted with a full set of event-year dummies ranging from one to four calendar years following drug approval (0 corresponds to FDA approval year).¹⁰ The coefficients on these interactions describe the corresponding proximity effect separately for each year.


The left panel of figure 3 plots how the estimated effect of proximity to a drug's first author on drug utilization evolves over time, while the right panel plots the effect of proximity to any of the drug's other pivotal clinical trial authors. The time pattern of proximity effects traced out in these graphs reveals a number of insights. First, recently approved cancer

⁸Data limitations prevent us from controlling for even finer divisions of cancer types. For each separate cancer type, we must observe at least two new drugs with different author locations to establish each region's counterfactual enthusiasm for new drug adoption within the diagnosis type. We chose to separate urologic cancers because they are typically treated by urologists who have completed a different type of fellowship training program (urology rather than hematology and oncology). It is common for hematologists to specialize in blood disorders, including leukemias and lymphomas, and not treat solid tumors. While oncologists may also subspecialize in treating particular types of solid tumors, their fellowship training spans cancer types, and it is not uncommon for the same therapeutic agent to find applications to more than one type of solid tumor. As a result, we group the solid tumors together.

⁹Both channels would be present even under random assignment of study author location.

¹⁰Medicare drug codes are not introduced until the calendar year following FDA approval for most drugs in our sample, limiting our ability to measure diffusion before the first year.

FIGURE 3.—INFLUENCE OF AUTHOR PROXIMITY ON DRUG USE



Graphs plot estimates of the effect pioneer investigator proximity has on drug utilization, t years since the corresponding cancer drug became FDA approved. Bands indicate 95% confidence intervals constructed from standard errors clustered at the provider HRR-drug level.

drugs are used more intensively on average in regions containing a study investigator, an effect that is much stronger in the first author's region. The second pattern highlighted by figure 3 is that the proximity effect fades over time, so that any proximity effect on drug utilization vanishes within four years after drug approval.

This figure provides a novel way to benchmark the speed of technology adoption. Prior measures of the speed of technology adoption have primarily focused on absolute rates of take-up, such as the length of time since invention for an individual to adopt (e.g., Comin & Hobijn, 2010) or rate of acceptance (e.g., Griliches, 1957). However, these measures can be inappropriate in settings where the "optimal" level of adoption is difficult to ascertain (e.g., due to informational uncertainty) and may even change over time, as competing technologies are introduced and scientific understanding evolves. In contrast, our measure of convergence describes how quickly regions conform to a benchmark adoption pattern set in regions containing the experts involved in the technology's development.

These estimates suggest that proximity to a pivotal trial investigator drives higher take-up of new drugs, an effect that is stronger and more persistent for first authors than for other authors of the pivotal study. Yet despite an initial eagerness to use the drug, this difference in diffusion between investigator and noninvestigator regions converges within a few years. This convergence provides further support for the econometric assumption that the drug's first author is not located in a geographic area with idiosyncratically higher latent demand for that particular cancer drug; if this were the case, we may

expect to see persistent differences in drug use across the first author and non-first author regions.

Table 3 shows results from our baseline specification in equation 1. Because the main proximity effects were found in figure 3 to be concentrated in the first two years following FDA approval, we focus the remainder of our regressions on this period. As shown in column 1 of table 3, we estimate that patients who receive cancer treatment in the first author's HRR are 4.0 percentage points more likely to receive the new drug, significant at the 1% level. To provide a useful benchmark, this first author impact is a 36% increase over the 11.1% average utilization rate in regions that contain a first author for a different in-sample drug with U.S.-based first authors. Patients who receive treatment in a middle or last author's region, by contrast, are only 0.69 percentage points more likely to receive the new drug, an estimate that is positive but not statistically significant. The difference between utilization in first author and other author HRRs is statistically significant at the 5% level.

Extent of author influence: Results by physician group. The average effect a study author has on prescribing behavior within the HRR may obscure important heterogeneity in regional utilization. The smaller average effect of "other" study authors relative to the first author could result from a narrower sphere of influence or from less enthusiastic adoption even within the other author's practice group. To explore this possibility, we use physician group tax IDs to measure which physicians practice in the same organization as a drug's trial authors.

TABLE 3.—AUTHOR PROXIMITY EFFECT ON DRUG UTILIZATION
Dependent Variable: New Drug Use

Proximity Measures	A. All HRRs		B. Author HRRs Only		C: New Cancer Patients	
	(1)	(2)	(3)	(4)	(5)	(6)
First author HRR	0.0404*** (0.0131)		0.0383*** (0.0122)		0.0399*** (0.0154)	
Other author HRR	0.0069 (0.0048)		0.0068 (0.0051)		0.0059 (0.0053)	
First author HRR and in author group		0.0421*** (0.0125)		0.0417** (0.0124)		0.0421*** (0.0162)
First author HRR and nonauthor group		0.0416** (0.0211)		0.0392** (0.0195)		0.0409** (0.0203)
Other author HRR and in author group		0.0276*** (0.0074)		0.0286*** (0.0073)		0.0271*** (0.0082)
Other author HRR and nonauthor group		-0.0031 (0.0054)		-0.0033 (0.0058)		-0.0036 (0.0057)
Number of observations	659,468	659,468	286,637	286,637	393,618	393,618

This table reports results from six separate regressions, where the dependent variable is an indicator that equals 1 if the patient is treated with the new drug over the observed episode of care. Each observation is a patient episode: a one-year episode of cancer treatment. The sample includes all episodes of care for the relevant cancer type within two years following a drug's FDA approval. "First author HRR" (hospital referral region) is an indicator variable that equals 1 if the patient is treated in the same HRR as the trial's first author. "Other author HRR" indicates treatment in the same HRR as another trial author (but not the first author). Author physician group indicators equal 1 if the patient was at any point during the year treated by a physician practicing in the same group as a trial author. All regressions include drug-year fixed effects; HRR-cancer-type fixed effects defined using three categories of cancer drugs: urologic, hematologic, and other (including breast, colon, lung, and brain); and indicators for patient age, race, sex, and new cancer treatment episode. Standard errors clustered at the HRR-drug level shown in parentheses. Significant at * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$. Columns 1 and 2: All HRRs and patient episodes within two calendar years after drug's FDA approval. Columns 3 and 4: Sample limited to HRRs containing any pivotal trial author. Columns 5 and 6: All HRRs but limited to new cancer patients, defined as patients with no cancer treatment in previous calendar year.

To estimate proximity effects separately within and outside an author's physician group, we estimate a modified version of equation (1) where the indicators for being treated in an author HRR are interacted with indicators for being treated by the author's physician group. Column 2 of table 3 shows the results of proximity separately by author group status. Among patients treated in the first author's HRR, patients treated within an author's physician group are 4.21 percentage points more likely to receive the new drug ($p < 0.01$) compared to patients treated in nonauthor regions, while patients treated outside the author group are 4.16 percentage points more likely to receive the drug. The point estimate suggests that the first author appears broadly influential, increasing new drug adoption in his or her region by almost equal amounts within and outside the practice group. However, the statistical significance of the utilization boost outside the first author's practice group is sensitive to the inference method, as discussed at greater length in the "Robustness" section that follows.

The results are quite different in other author regions. In those regions, patients treated within an author's practice group are 2.8 percentage points more likely to receive the new drug ($p < 0.01$). However, there is no estimated increase in drug utilization outside the author group, and the 95% confidence interval is bounded above by 0.7 percentage points. Despite being enthusiastic adopters of new drugs within their own practice group, middle and last authors do not appear to influence practice patterns in neighboring physician groups. (See appendix table A2 for a version of table 3 that separates middle and last authors.)

An important consideration for interpreting these results is that because patients are not randomly assigned to doctors, some of the increased utilization found within author group could be driven by a compositional shift in which patients are treated by author group physicians. To the extent

that patients most appropriate for the new drug sort into the author group, some of the increase in prescribing rates within the physician group may not correspond to a net increase in propensity to treat a given patient with the new drug. The region level results in column 1, however, indicate that overall prescribing does in fact increase at the region level.

Taken as a whole, the results from columns 1 and 2 provide evidence of important proximity effects within author regions. The higher rates of drug utilization in a first author region compared to other study author regions appear to be driven primarily by the first author's broader sphere influence. While both types of study authors boost utilization within their own practice group, the evidence suggests that only first authors substantially increase drug adoption by doctors who are not part of the author's group.

Robustness. To probe the robustness of the baseline proximity effects reported in panel A of table 3, we run an identical analysis in a restricted sample that includes only regions with at least one study author for a drug in our sample. Restricting the sample mitigates the concern that non-research-intensive regions provide a poor counterfactual for a new drug's popularity in author regions. Results are reported in table 3, panel B, with estimates very similar to those found in the full sample reported in panel A. We continue to find approximately 4 percentage point higher new drug use in the first author region ($p < 0.01$) and a smaller, insignificant aggregate effect on other author regions. Notably, restricting the sample in this way does not substantially attenuate the estimated effects.

The interpretation we emphasize for our baseline results is that increased utilization in author HRRs results from knowledge spillovers from proximity to lead investigators. However, a competing explanation is that heightened levels

of utilization among patients in author HRRs might simply reflect the continuation of treatment for patients enrolled in the trial itself. A direct test of the continuation hypothesis would be to identify patients in our data who were enrolled in the pivotal trial and remove them from the analysis. While the Medicare data do not permit identifying patients enrolled in specific clinical trials, we reestimate our baseline results in a sample restricted to new cancer patients, defined to be those with no cancer treatment observed in the previous calendar year. This sample restriction plausibly removes patients enrolled in the pivotal study and eligible for Medicare who survive and continue to receive the new drug following FDA approval. Estimates based on new cancer patients alone, reported in panel C of table 3, are nearly identical to our baseline estimates in panel A of table 3 that include all cancer patients. We find these results to support the interpretation that our baseline effects are driven by informational spillovers rather than a continuation mechanism alone. Further, a back-of-the-envelope calculation described in appendix A suggests that our effect size is too large to be explained by the number of patients likely enrolled in the author's clinical trial site.

Finally, we explore the robustness of our findings to bootstrap methods of statistical inference, with results reported in appendix table A3. Our main results use conventional cluster robust variance estimators, supported by the fact that we have many clusters (6,086 region-drug pairs) of which 18 clusters are "treated" (i.e., first author region for given drug). To account for possible bias due to a modest number of unequally sized treated clusters, we apply the wild cluster bootstrap developed by Cameron, Gelbach, and Miller (2008) and Cameron and Miller (2015). We also vary the cluster definition, estimating p -values with clusters at the HRR-drug level as in our baseline results, as well as at the broader HRR level. When the number of treated clusters is small, results from MacKinnon and Webb (2016) suggest that conventional clustered standard errors are likely to overreject, but the wild cluster bootstrap tends to underreject. As a result, the wild cluster bootstrap estimates should be conservative, with the true p -values lying between the values given by each approach.

Our main findings remain statistically significant across all four methods of inference reported in appendix table A3. In particular, we consistently find that first author regions have significantly higher drug use than nonauthor regions (p -value ranges from 0.002 to 0.012, depending on clustering method). We also find that the impact of being treated by the author's group in a region that contains a middle or last author consistently significant across inference approaches ($p \leq 0.002$ with all methods).

One result is sensitive to the inference method: our finding that patients treated in the first author's region but outside the first author's group are significantly more likely to get the new drug than patients in nonauthor regions. The p -value is marginally significant with conventional cluster robust standard errors ($p = 0.048$) and becomes not

statistically distinguishable from 0 with the wild cluster bootstrap ($p = 0.168$ or 0.278 depending on cluster level). While the point estimates for utilization in the first author region within and outside the author's practice group are very close, we caution that the estimate of utilization rates outside the author's group is imprecisely measured.

Mechanisms of estimated author influence. The greater impact of proximity to first authors could be driven by two potential factors: (a) first authors take on more responsibility for analyzing and writing the paper, and thus are better informed about the new drug's value; or (b) even if all authors had the same quality of information about a new drug, first authors may be more influential due to their greater professional stature. Both channels have potential a priori support.

The pivotal clinical trials for drugs in our study have an average of 14 authors per paper (table 1). In the publication of these trials, the first author often takes the lead role in trial design and preparing the manuscript (Hudes et al., 2007; DeAngelo et al., 2007), suggesting he or she may also have the most detailed, comprehensive view of the drug's efficacy. The first author is also likely to be one of the highest-profile physicians involved with the research; he or she is the single most highly cited clinical author for eight of the seventeen drugs in our sample with U.S.-based trials. The findings reported in the previous section on authors' scope of influence outside their practice group lends support to the idea that the first author's status as a local opinion leader may primarily drive these differences. In the next section, we further explore whether identifying superstar authors based on citation histories rather than authorship sequence leads to similar findings.

Another complementary explanation for the observed impact of first authors is that they may have stronger ties to the sponsoring drug companies and be more actively involved in drug promotion efforts. Of the 21 drugs in our sample, 9 of the published trials report disclosure statements detailing which authors have financial relationships with drug companies. For these drugs, an average of 52% of all clinical authors report financial ties to the sponsoring drug company compared to 67% of first authors. These financial ties include consulting fees, lecture fees, research support, expert testimony, and stock ownership.

While drug companies are only 1.3 times more likely to have financial ties to a first author compared to a middle or last author, the estimated impact of being treated in the first author's region on new drug use is over 5.5 times larger than the estimated impact of being treated in another author's region. If each disclosed financial tie indicated an equal amount of funding support and financial ties were the only driver of our observed effect, then we would expect differences in the frequency of financial relationships to scale linearly with the estimated effect of author proximity. Hence, the observed frequency of financial ties between drug companies and clinical authors would not lead us to predict the

first author’s apparent outsized influence on regional drug utilization.

Unfortunately, our assessment is limited by the fact that many drug trials did not report disclosures over this period, and trials with disclosures did not list the amounts of money exchanged. Thus, we cannot rule out the possibility that drug companies have stronger relationships or expend more resources supporting the first author’s drug promotion efforts compared to other authors. Crucially, if drug companies were investing more in the first author, this would suggest that they perceived a higher return to the first author’s potential promotion efforts; in that case, the drug company’s investment is complementary to the superior information or professional stature that the first author already offers.

A related issue is that because the geographic location of a study author is also typically the location of a trial site (though one of many), it is difficult to disentangle whether local physicians primarily learn about a new drug because of proximity to a study author or proximity to trial activity. However, our finding of a larger first author proximity effect compared to other authors sheds some light on this issue. Specifically, the differential effects do not appear to be easily explained by substantial differences in the size of the trial site by author type: for the eleven drugs in our study for which we could determine the author location of the largest trial site in the study, eight (73%) were not located in the lead author’s HRR.

Outside of drug company–sponsored events, there are many other opportunities for oncologists to meet with their local peers and share ideas. Within a given oncology practice group, formal mechanisms may include the establishment of internal drug treatment protocols and “tumor board” meetings where treatment options for new cancer cases are often discussed with a broad team of care providers. Across separate practice groups, opportunities for sharing ideas include invited grand rounds seminars, local and regional professional society meetings, contact through shared patients and patient referrals, and casual networking.

In personal communications with oncologists, physicians described significant barriers to the adoption of new cancer agents. Because trial participants are often selected for being in more stable health than many cancer patients, physicians cited significant concern about the risk of severe side effects and uncertainty about optimal dosing regimens. Oncologists are also aware of the potential for heterogeneous treatment response; hearing about successfully treated patients may be more compelling than reading about modest average response rates. The expertise of a prominent physician in the community on a new drug’s applications and efficacy could substantially lessen these barriers to new drug adoption.

Finally, another explanation for the observed regional differences in drug utilization is differences in drug prices. There are two relevant prices to consider in this context: the reimbursement that physicians are paid for prescribing a drug and the cost to the physician to purchase the drug. If drug-specific reimbursements were higher or purchasing

costs lower in first author regions than in other regions, prices could potentially explain our proximity findings. However, we think neither of these price effects is a likely explanation in our context. Because we are studying utilization among Medicare patients where reimbursement rates are set by administrative rule, there is effectively no scope for reimbursements to be idiosyncratically higher in a given region for first author drugs. In addition, while it is conceivable that drug manufacturers could offer drug discounts to specific physicians or groups and might specifically target physicians practicing at the first author’s hospital, our high estimated utilization rate for local physicians who do not practice in the first author’s physician group makes it seem unlikely that drug discounts could explain our results.

C. Superstar Proximity Effects

Our baseline results in section IIIB suggest that clinical trial authors with greater expertise or prominence (as captured by first author status) have a greater impact on drug utilization in their region. To further explore the differential effect of superstar physicians, we apply our citation-based measure of author prominence that identifies the top 10% (or 50%) authors as the top-cited author on that particular drug trial, plus any other authors whose citation count places them in the top 10% (50%) of all authors on the same drug’s trial. Because we want to make comparisons of doctor influence within the set of physicians on a particular drug trial, this relative measure provides a clear comparison that will ensure our ability to separate the relative prominence of authors on each trial. Note that because trial authors are a selected population, it is likely that these citation ranks would be even more favorable if compared to the overall population of physicians in their field.

For any measure of superstar status, our baseline regression in equation 1 is easily modified to estimate the differential impact of proximity to a superstar author. For these regressions, we allow author proximity effects to vary by drug and then estimate the differential impact of proximity to a superstar author. Our superstar regression takes the form

$$\begin{aligned} (\text{drug})_{ijt} = & \beta_s \mathbf{1}(\text{superstar author HRR})_{jd} \\ & + \{\text{author HRR} \times \text{drug FEs}\}_{jd} \\ & + \{\text{HRR} \times \text{disease-group FEs}\}_{ijt} \\ & + \{\text{drug} \times \text{year FEs}\}_{dt} + \delta X_{it} + \varepsilon_{ijt}. \end{aligned} \quad (2)$$

The key coefficient of interest is β_s , which describes how much more a new drug is used in a superstar HRR relative to other author HRRs, on average. Thus, if $\beta_s = 0$, utilization in a superstar region is no more intensive than in other author HRRs, while $\beta_s > 0$ corresponds to higher utilization in superstar regions. The second term in this regression allows the effect of author proximity to vary by drug, and the last three terms are the same as in equation 1.

The results from regression 2 are shown in table 4. Column 1 shows that drug utilization in first author HRRs is

TABLE 4.—SUPERSTAR AUTHOR PROXIMITY EFFECT ON DRUG UTILIZATION
Dependent Variable: New Drug Use

Independent Variables	(1)	(2)	(3)	(4)	(5)	(6)
First author HRR	0.0300** (0.0124)			0.0246** (0.0119)	0.0242** (0.0111)	0.0230** (0.0110)
Top 50% cited author HRR		0.0227** (0.0090)		0.0154* (0.0081)		0.0144 (0.0090)
Top 10% cited author HRR			0.0232** (0.0122)		0.0100 (0.0108)	0.0034 (0.0117)
Number of observations	659,468	659,468	286,637	286,637	393,618	393,618

This table reports results from six regressions that test whether superstar authors are more influential than other study authors for the same drug. The baseline regression specification is augmented to include a vector of (drug)×(any author HRR) fixed effects. Reported coefficients describe whether regions with authors of the noted type have higher new drug use compared to the rest of the author regions for the same drug. Top 50% and top 10% authors are defined as the most prominent academic authors for each drug, as measured by citation counts accruing to publications produced over the ten years leading up to FDA drug approval in the relevant field. See notes to table 3. * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

3.00 percentage points higher on average than in middle or last author HRRs for the same drug, which closely matches the result obtained by differencing the first and other author HRR results in table 3. From columns 2 and 3, new drug utilization is 2.27 percentage points higher in a top 50% cited author HRR (column 2) and 2.32 percentage points higher in a top 10% cited author HRR.

Columns 4 to 6 run horse races between these three superstar measures. Column 4, which includes superstar indicators for both first author, and top 50% cited author, shows that both indicators correspond to higher utilization; the coefficient on the top 50% cited author is marginally significant, with $p = 0.056$. Column 5 includes first author and top 10% cited author indicators; both coefficients point to higher utilization, although only the first author corresponds to a statistically significant increase. Similarly, in column 6, which includes all three superstar measures, all coefficients are positive, but only first author is significant.

Each of the author regions contains a physician investigator who is well informed about the new drug, but among author regions, those with the most prominent authors are the ones that experience the most substantial increases in new drug use. These findings suggest an important role of local opinion leaders even in the context of drug adoption by highly expert decision makers with access to clinical trial findings. Regional information frictions may dissipate within three to four years, but during the initial two years after drug introduction, local opinion leaders have substantial influence on adoption rates in their region. These results do not estimate the causal impact of increasing an individual author's citation history or authorship order *ceteris paribus*; principal investigators and authors with high citation counts are likely to be exceptional along other unmeasured dimensions as well.

D. Extent of Investigator Influence

Appendix table A4 further probes the reach of drug authors' influence. First, we investigate whether the impact of investigator proximity is related to regional enthusiasm for other new cancer drugs. We find that the first author's influence is greatest in regions that are typically slower to adopt new drugs. There may be greater scope for the study

author to affect practice patterns in slower-adopting regions that are not already very high users of new cancer drugs.

Second, we test whether study authors affect new drug utilization in neighboring regions. Although the first author's influence may extend beyond physicians in his or her own practice group to others practicing in the same region, there is no evidence that his or her influence raises new drug utilization in neighboring regions.

Finally, we study author influence on off-label drug use. We find no evidence of higher use of the drug for off-label patients in the authors' regions, suggesting authors' influence is largely local to the cancer type on the initial label. Notably, the estimates on off-label drug use are less precise than our baseline findings, given the lower rate of off-label prescribing.

IV. Patient Travel and Selective Sorting

As discussed in section IIIA, there are two possible channels through which the observed increased propensity to prescribe new drugs in first author regions may occur: an increased propensity to use the drug on a fixed set of patients and a change in patient sorting such that the first author regions see patients with higher latent demand. In this section, we test directly for changes in patient sorting as indicated by patient travel patterns and then use an instrumental variables strategy to identify the differences in drug utilization that occur over a fixed set of patients.

In table 5, we begin by testing whether patients with the targeted diagnosis who seek treatment in the first author's HRR are more likely to have traveled from a different HRR of residence. This would occur if, for example, savvy patients travel into author regions for treatment in order to gain access to the new cancer drugs. In columns 1 and 2, the regression specification mirrors that in the main specification described in equation (1), but the outcome variable has been replaced with an indicator variable for travel, defined by whether the patient's HRR of residence is not the same as the HRR where he or she receives care.

In the baseline specification reported in column 1, we find a 3.3 percentage point increase in the fraction of patients treated in the first author's HRR who do not live in the region, significant at the 10% level; on average, 22.2% of patients

TABLE 5.—PATIENT TRAVEL AND PROXIMITY EFFECTS

Independent Variables	Dependent Variables			
	Travel		New Drug Use	
	(1)	(2)	(3)	(4)
First author HRR	0.0329* (0.0192)	0.0309* (0.0197)	0.0327*** (0.0124)	0.0311*** (0.0116)
Traveler to first author HRR			0.0224 (0.0140)	0.0226 (0.0138)
Other author HRR	0.0295*** (0.0094)	0.0285*** (0.0089)	0.0066 (0.0052)	0.0064 (0.0056)
Traveler to other author HRR			0.0012 (0.0061)	0.0011 (0.0064)
Sample				
Author HRRs only?	No	Yes	No	Yes
Number of observations	659,468	286,637	659,468	286,637

Columns 1 and 2 report results from regressions where the dependent variable indicates whether the patient received care outside his or her HRR of residence. Columns 3 and 4 report results from regressions where the dependent variable indicates whether the patient received treatment with the new drug. Columns 2 and 4 restrict the sample to regions that contain a study author for at least one drug in our sample. See notes to table 3. * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

treated in an author region reside outside the region. There is a similar 3.0 percentage point increase in the fraction of traveling patients treated in other authors' HRRs, significant at the 1% level. Restricting to the set of regions that ever contain a study author in column 2 yields similar results.

This evidence suggests that some patients are aware of new centers of expertise for the new cancer drug (perhaps due to physician referral) and are willing to travel farther to improve their access to the drug. Note that these findings do not necessarily require patients to cross large distances for care; 66% of travelers are being treated in a neighboring region that shares a border with the patient's HRR of residence. In an unreported regression, we found that these travel effects are driven by patients seeking care in an author's practice group; excluding patients treated by an author group leads to negative, insignificant coefficients in the travel regression. Taken together, these findings are consistent with the possibility that some patients are on the margin between seeking care at their local tertiary care center or traveling to a neighboring, prestigious academic medical center for treatment. If the center has recently participated in a clinical trial for a new drug treating the patient's cancer, the patient may be more likely to travel for treatment.

If these patients who are newly traveling into an author's practice group are either more clinically appropriate for the new drug or have higher demand for trying the new technology, then part of the increased levels of drug utilization in the author's region may be driven by the changing patient composition. In table 5, columns 3 and 4, we test whether the patients who travel from outside HRRs differ in their propensity to receive the new drug relative to nonmovers. In these columns, we report results from a regression that augments our baseline equation (1) by interacting the author proximity indicators with a binary indicator *traveler* for whether the observed patient is seeking care outside his or her HRR of residence. Based on this regression, we estimate that patients traveling to the first author's HRR are 2.2 percentage points more likely to receive treatment with the new

drug than patients treated in the first author HRR who also reside within that HRR, although the result is not statistically significant at conventional levels. Travelers to other author regions are only 0.1 percentage points more likely to get the drug, which is also not significant. The confidence intervals on these travel estimates allow the possibility that the overall 4.0 percentage point higher new drug use in first author regions may be driven at least in part by changing patient composition, and not solely by a higher propensity to use the drug on a fixed set of patients.

These estimates on patient travel are particularly relevant considering the new attention to provider networks available on the Affordable Care Act's health insurance exchanges. A common feature of these new insurance plans is restricted provider networks (Hancock, 2013), with consumers facing much higher prices for out-of-network care. Our findings on travel suggest that severely ill patients, such as the cancer patients in our study, may travel strategically to improve access to providers with additional expertise in new treatments.

Instrumental variable analysis. To isolate whether trial author regions are indeed more likely than other regions to use the drug on a given set of patients, we pursue an instrumental variables (IV) strategy. In particular, we use indicators for whether the patient resides in the first or other author's HRR as an instrumental variable to predict whether he or she seek treatment in an HRR that contains the first or other author for the relevant drug. This instrumentation strategy mitigates the concern that patient sorting renders the patients treated in the first author region more suitable to treatment with the new cancer drug.

The reduced-form equation of the IV model takes the following form:

$$\begin{aligned}
 (\text{drug})_{ijd} = & \gamma_1 \mathbf{1}(\text{reside in first author HRR})_{ijd} \\
 & + \gamma_2 \mathbf{1}(\text{reside in other author HRR})_{ijd} \\
 & + \{\text{HRR} \times \text{disease-group FEs}\}_{ijd} \\
 & + \{\text{drug} \times \text{year FEs}\}_{dt} + \delta X_{it} + \varepsilon_{ijd}. \quad (3)
 \end{aligned}$$

Paralleling the baseline regression specification, we include fixed effects for HRR by disease group and for drug by year. We also report results from an enriched IV specification, where in addition to using the two indicators for residence in an author HRR as instrumental variables, we also include two additional instruments: (a) residence in a first author's neighboring HRR (a region that shares a border with the first author HRR) and (b) residence in an other author's neighboring HRR (a region bordering an other author HRR).

The IV exclusion restriction requires that after conditioning on the included fixed effects, where a patient lives is uncorrelated with his or her suitability or demand for treatment with the new cancer drug. For example, because we include region by drug class fixed effects, this allows regions to vary in their latent demand for drug classes but does not allow author regions to have higher latent demand for the

TABLE 6.—IV ESTIMATES OF PROXIMITY EFFECT ON DRUG UTILIZATION

	Outcome: New Drug Use		
	(1)	(2)	(3)
A. Reduced form: Drug receipt effect			
Residence in first author HRR	0.0228** (0.0099)	0.0217** (0.0099)	0.0228** (0.0099)
Residence in other author HRR	0.0038 (0.0044)	0.0044 (0.0050)	0.0046 (0.0061)
Residence in first author's neighbor HRR			0.0035 (0.0045)
Residence in other author's neighbor HRR			−0.0028 (0.0028)
B. Two-stage least squares			
Provider in first author HRR	0.0293** (0.0127)	0.0259** (0.0114)	0.0263** (0.0115)
Provider in other author HRR	0.0056 (0.0058)	0.0060 (0.0061)	0.0065 (0.0061)
Sample			
Author HRR only?	No	Yes	No
Number of observations	659,468	286,637	659,468

Reduced-form results report coefficients from three regressions where the outcome variable is new drug use and the key explanatory variables are indicators for whether a given patient resides in the same region as the study author (or in column 3, in a neighboring region). In panel B, the two-stage least squares results use patient residence variables as instrumental variables for whether the patient is treated in the author region. All regressions include drug-year fixed effects and HRR-cancer type fixed effects defined using three categories of cancer drugs: urologic, hematologic, and other (including breast, colon, lung, and brain). Standard errors clustered at the HRR-drug level shown in parentheses. * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

author's drug compared to other drugs in the same class. The exclusion restriction could be violated under a few conditions. One possibility is that patients with the targeted cancer who reside in the first author region could have idiosyncratically high demand for the drug; this could occur if, for example, the drug targets a particular subtype of colon cancer that has a higher-than-typical prevalence in the first author's region, so that a larger fraction of colon cancer patients in the region is appropriate for treatment. Second, the instrument would be invalid if patients change their HRR of residence in response to the availability of new cancer drugs.

While the IV exclusion restriction is not directly testable, it seems plausible that the fraction of targeted cancer patients suitable for treatment with the new drug would not vary systematically across regions and that elderly Medicare patients would be very unlikely to move across regions within a three-year period in response to the location of a new cancer drug trial. This assumption is bolstered by the observed convergence in drug usage across first author and non-first author regions, as reported in figure 3, suggesting no permanent differences in patient eligibility for treatment in the first author's region.

Results from the IV regressions are reported in table 6. The reduced-form results show that patients residing in the first author's HRR are 2.3 percentage points more likely to receive treatment with the new drug, significant at the 5% level; there is no significant increase in use for patients residing in other author regions. The IV estimate reported in the final rows of the table rescales the reduced-form estimate and shows that providers in the first author's region are 2.9 percentage points more likely to prescribe the new drug compared to other providers, significant at the 5% level. The finding is robust to restricting the sample to patients residing

in HRRs that contain an author for any drug, as reported in column 2. Adding the instrumental variables for residence in neighbor HRRs to the model also does not substantially change the estimated IV coefficient (cf. column 3).

The IV results suggest that over 70% of the baseline effect reported in section IIIB is due to the increased propensity of physicians in the first author's region to prescribe the drug on a given set of patients. While the differential sorting of high-demand or high-appropriateness patients to author regions explains some of the observed boost in drug utilization, most of the effect is driven by differences in physician behavior, not patient sorting. Under a local average treatment effect interpretation, the IV result implies that doctors in the first author's region are 2.9 percentage points more likely to use the new drug on a given set of patients—those for whom location of residence determines location of care.

Taken together, the patient traveling results and the IV regressions find support for both hypothesized channels by which the presence of a first author may affect care in his or her region. Patients with high latent demand for the drug seem to seek out care in areas with high expertise in the new technology. In addition, doctors in the first author's region are more likely to use the new drug, holding fixed the population of patients seeking treatment, with this channel representing the primary driver behind our baseline result.

Extrapolating from the IV regression result, if all providers behaved like those in the first author's region, approximately 2,500 additional Medicare fee-for-service (FFS) patients would be treated with each new cancer drug in the first two years after initial drug approval. This amounts to an estimated 53,000 Medicare FFS patients in total over the eleven years of our sample who did not receive treatment with one of the 21 drugs under study due to the lower patterns of initial usage in areas that did not contain the study's first author.

V. Welfare Implications of Early Drug Diffusion

The evidence presented thus far suggests that clinical trial investigators influence local drug use in the first few years following drug introduction. A key remaining question is whether faster drug diffusion into late-adopting regions would be welfare enhancing. Are nonauthor regions slower to adopt new drugs because they are realizing more modest survival benefits?

Contrasting the survival benefits to new drug adoption in author and nonauthor regions is not statistically powered in our setting, given the modest number of author regions and the limited impact of most new drugs on survival. If we maintain the identification strategy specified in equation (1) but replace the outcome variable with one-year mortality, we estimate that one-year mortality is 0.65 percentage points lower in first author regions, with the 95% confidence interval running from −2.6 to 1.3 percentage points. The corresponding Wald instrumental variable estimate for the effect

of a new drug is similarly imprecise, suggesting a statistically insignificant 15 percentage point reduction in one-year mortality.

Due to the lack of statistical power for this analysis, we turn to a broader comparison of cancer mortality in fast- and slow-adopting regions to identify whether the survival gains from new cancer drugs are related to observed adoption patterns in our data. (Details of the difference-in-differences estimation are described in appendix figures A1 and A2, as well as accompanying notes.) The key source of identification comes from comparing survival rates across fast- and slow-diffusion regions before and after the introduction of a new cancer drug. For this strategy, we measure a region's enthusiasm for adopting a new drug using a leave-out estimate of enthusiasm for other new drugs, excluding the region's realized utilization for that particular new drug.

We find that new drug use is 3.0 percentage points higher in fast-diffusing regions in the first four years following FDA approval ($p < 0.01$). Further, regions that are fast adopters experience a 1 percentage point greater reduction in one-year mortality for the relevant cancer diagnosis than slow-adopting regions ($p < 0.05$). The effect is even larger for eighteen-month mortality, with fast adopters experiencing a 1.5 percentage point greater reduction than slow adopters ($p < 0.01$). If we interpret the effect on drug utilization as the first stage of an instrumental variables regression to calculate the local average treatment effect of new drug use, the Wald instrumental variable estimate would suggest that new drug use is associated with a 51 percentage point reduction in eighteen-month mortality, among the marginal treated patients in high-diffusing regions, with a 95% confidence interval of $[-0.83, -0.19]$.

However, the implied instrumental variable estimate of the drug treatment effect stretches credibility, considering that the average eighteen-month mortality rate is 35.8% in our sample. It could be that the marginal patient who would be treated in a fast-adopting region but not in a slow-adopting region has much higher-than-average latent mortality. However, a more compelling explanation may be that fast-diffusion regions have higher returns to new drug use than slow-diffusion regions even for inframarginal patients. Due to superior dosing, design of drug cocktails, management of side effects, or patient selection, these regions may achieve higher average returns to new drug use among all treated patients.¹¹

Considered in the context of our earlier results, these findings suggest that policies designed to increase drug utilization without improving the surrounding information environment on appropriate application and management of

the new drug may not achieve substantial survival gains. Information frictions may be a key limiting factor not only for drug adoption but also for the medical benefits of new drug use. Interventions that improve local information about a new drug, such as through closer contact with experts, may result in higher use and improved patient welfare by facilitating the effective application and targeting of the new technology.

VI. Conclusion

The results suggest that information frictions significantly limit the adoption of new cancer drugs in the first few years after drug introduction. Prominent physicians who are well informed about a new drug in the early stages of diffusion may play a key role in easing these frictions. Our results show that new drug utilization is 36% higher over the two years following FDA approval in the hospital market where the lead physician investigator on the drug's pivotal clinical trial practices. Despite the marked regional differences in early adoption of new cancer drugs, there is no evidence that early expertise with a drug drives higher rates of long-term utilization. Author HRRs are no more likely to specialize in treatment with the new drug than other regions by the fourth year following drug introduction. Thus, the information frictions that may hamper early adoption seem to ease over time as utilization rates in first author and other HRRs converge within a four-year period.

While both first authors and other authors increase new drug utilization within their physician practice group, the influence of first authors is larger. Further, we show that relying on publication and citation history rather than authorship order to identify superstar study authors yields similar results, suggesting that investigators with greater professional status are more influential in increasing drug adoption in their region. We also find that investigator proximity has the largest impact on regions with the lowest levels of adoption of other new drugs, pointing to greater adoption frictions in less technology-intensive areas.

Regions that are quick to adopt new cancer drugs also see greater improvements in patient survival following new drug introduction. The large magnitude of these survival gains in fast-adopting regions suggests they may be driven not only by the greater fraction of patients receiving treatment with the new drug in these regions but also by higher average returns to new drug use.

This study has important limitations that may point the way to informative avenues for future research. While the results suggest a complementarity between the opinion leader's professional prominence and superior information about a new drug, we cannot fully disentangle how information and prominence each contributes to local diffusion. Future work should assess the role of local opinion leaders outside the context of research involvement to better understand the potential policy mechanisms that could improve local dissemination of information about medical innovations. Finally, a more detailed understanding of the factors

¹¹ Oncologists have documented dramatic variation in the return to drug use corresponding to variation in information regarding appropriate use of the drug. For example, when trastuzumab was introduced, breast cancer patients taking the drug survived for a median of 25 months. Now, median survival has increased to 41 months for patients treated with the same drug, which expert oncologists have attributed to longer dosing schedules and better side effect management (Pollack, 2014).

shaping the heterogeneous estimated returns to new technology adoption will be important for understanding the welfare consequences of variation in early diffusion of medical treatment.

The evidence in this paper suggests that the local information environment may be a key determinant of both the adoption of new cancer drugs and the realized survival benefits. Merely increasing new drug use in slower-adopting regions may not lead to substantial improvements in patient survival, but improving the local information environment to support new drug adoption could have significant welfare benefits by increasing appropriate utilization and improving survival rates for treated patients. Contact with prominent physicians with expertise on a new drug may increase local drug use by reducing information frictions.

REFERENCES

- Adhvaryu, Achyuta, "Learning, Misallocation, and Technology Adoption: Evidence from New Malaria Therapy in Tanzania," *Review of Economic Studies* 81 (2014), 1331–1365.
- Audretsch, David B., and Maryann P. Feldman, "Knowledge Spillovers and the Geography of Innovation" (vol. 4, 2713–2739), in J. Vernon Henderson and Jacques-François Thisse, *Handbook of Regional and Urban Economics* (Amsterdam: Elsevier, 2004).
- Azoulay, Pierre, Joshua S. Graff Zivin, and Bhaven N. Sampat, *The Diffusion of Scientific Knowledge across Time and Space: Evidence from Professional Transitions for the Superstars of Medicine* (Chicago: University of Chicago Press, 2012).
- Bach, Peter B., "Limits on Medicare's Ability to Control Rising Spending on Cancer Drugs," *New England Journal of Medicine* 360 (2009), 626–633.
- Baerlocher, Mark Otto, Marshall Newton, Tina Gautam, George Tomlinson, and Allan S. Detsky, "The Meaning of Author Order in Medical Research," *Journal of Investigative Medicine* 55 (2007), 174.
- Baicker, Katherine, and Amitabh Chandra, "Understanding Agglomerations in Health Care" (pp. 211–236), in Edward L. Glaeser, ed., *Agglomeration Economics* (Chicago: University of Chicago Press, 2010).
- Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson, "The Diffusion of Microfinance," *Science* 341 (2013).
- Bristow, Robert E., Jenny Chang, Argyrios Ziogas, Hoda Anton-Culver, and Veronica M. Vieira, "Spatial Analysis of Adherence to Treatment Guidelines for Advanced-Stage Ovarian Cancer and the Impact of Race and Socioeconomic Status," *Gynecologic oncology* 134 (2014), 60–67.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *this REVIEW* 90 (2008), 414–427.
- Cameron, A. Colin, and Douglas L. Miller, "A Practitioner's Guide to Cluster Robust Inference," *Journal of Human Resources* 50 (2015), 317–372.
- Chandra, Amitabh, and Douglas O. Staiger, "Productivity Spillovers in Health Care: Evidence from the Treatment of Heart Attacks," *Journal of Political Economy* 115 (2007), 103–140.
- Coleman, James, Elihu Katz, and Herbert Menzel, "The Diffusion of an Innovation among Physicians," *Sociometry* 20 (1957), 253–270.
- Comin, Diego A., Mikhail Dmitriev, and Esteban Rossi-Hansberg, "The Spatial Diffusion of Technology," NBER working paper 18534 (2012).
- Comin, Diego, and Bart Hobbijn, "An Exploration of Technology Diffusion," *American Economic Review* 100 (2010), 2031–2059.
- Conley, Timothy G., and Christopher R. Udry, "Learning about a New Technology: Pineapple in Ghana," *American Economic Review* 100 (2010), 35–69.
- Crespi, Gustavo, Chiara Criscuolo, Jonathan E. Haskel, and Matthew Slaughter. 2008. "Productivity Growth, Knowledge Flows, and Spillovers," Working Paper 13959, National Bureau of Economic Research.
- DeAngelo, Daniel J., Daohai Yu, Jeffrey L. Johnson, Steven E. Coutre, Richard M. Stone, Alison T. Stopeck, Jon P. Gockerman, Beverly S. Mitchell, Frederick R. Appelbaum, and Richard A. Larson, "Nelarabine Induces Complete Remissions in Adults with Relapsed or Refractory T-Lineage Acute Lymphoblastic Leukemia or Lymphoblastic Lymphoma: Cancer and Leukemia Group B Study 19801," *Blood* 109 (2007), 5136–5142.
- Escarce, Jose J., "Externalities in Hospitals and Physician Adoption of a New Surgical Technology: An Exploratory Analysis," *Journal of Health Economics* 15 (1996), 715–734.
- Fairfield, Kathleen M., F. Lee Lucas, Craig C. Earle, Laurie Small, Edward L. Trimble, and Joan L. Warren, "Regional Variation in Cancer-Directed Surgery and Mortality Among Women with Epithelial Ovarian Cancer in the Medicare Population," *Cancer* 116 (2010), 4840–4848.
- Farrow, Diana C., William C. Hunt, and Jonathan M. Samet, "Geographic Variation in the Treatment of Localized Breast Cancer," *New England Journal of Medicine* 326 (1992), 1097–1101.
- Fisher, Elliott S., David E. Wennberg, Thérèse A. Stukel, Daniel J. Gottlieb, F. L. Lucas, and Étoile L. Pinder, "The Implications of Regional Variations in Medicare Spending. Part 1: The Content, Quality, and Accessibility of Care," *Annals of Internal Medicine* 138 (2003a), 273–287.
- "The Implications of Regional Variations in Medicare Spending. Part 2: Health Outcomes and Satisfaction with Care," *Annals of Internal Medicine* 138 (2003b), 288–298.
- Griliches, Zvi, "Hybrid Corn: An Exploration in the Economics of Technological Change," *Econometrica* 25 (1957), 501–522.
- Hancock, Jay, "Narrow Networks' Trigger Push-Back from State Officials," Kaiser Health News. <http://www.kaiserhealthnews.org/stories/2013/november/25/states-balk-at-narrow-networks.aspx>.
- Hudes, Gary, Michael Carducci, Piotr Tomczak, Janice Dutcher, Robert Figlin, Anil Kapoor, Elzbieta Staroslawska, et al., "Temsirolimus, Interferon Alfa, or Both for Advanced Renal-Cell Carcinoma," *New England Journal of Medicine* 356 (2007), 2271–2281.
- Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson, "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations," *Quarterly Journal of Economics* 108 (1993), 577–598.
- MacKinnon, James G., and Matthew D. Webb, "Wild Bootstrap Inference for Wildly Different Cluster Sizes," *Journal of Applied Econometrics* 32 (2016), 233–245.
- Marshall, Alfred, *Principles of Economics* (London: MacMillan, 1890).
- Nattinger, Ann Butler, Mark S. Gottlieb, Judith Veum, David Yahnke, and James S. Goodwin, "Geographic Variation in the Use of Breast-Conserving Treatment for Breast Cancer," *New England Journal of Medicine* 326 (1992), 1102–1107.
- Newhouse, Joseph P., "Medical Care Costs: How Much Welfare Loss?" *Journal of Economic Perspectives* 6 (1992), 3–21.
- Onega, Tracy, Eric J. Duell, Sun Shi, Dongmei Wang, Eugene Demidenko, and David Goodman, "Geographic Access to Cancer Care in the US," *Cancer* 112 (2008), 909–918.
- Phelps, Charles E., "Information Diffusion and Best Practice Adoption" (vol. 1, pp. 223–264), in Anthony, J. Culyer and Joseph P. Newhouse, eds., *Handbook of Health Economics* (Amsterdam: Elsevier, 2000).
- Pollack, Andrew, "Roche Breast Cancer Drug Perjeta Appears to Greatly Extend Patients' Lives," *New York Times*, September 28, 2014.
- Rogers, Everett M., *Diffusion of Innovations* (New York: Free Press, 1962).
- Skinner, Jonathan, and Douglas Staiger, "Technology Adoption from Hybrid Corn to Beta Blockers," NBER working paper 11251 (2005).
- Smith, Sheila, Joseph P. Newhouse, and Mark S. Freeland, "Income, Insurance, and Technology: Why Does Health Spending Outpace Economic Growth?" *Health Affairs* 28 (2009), 1276–1284.
- U.S. CBO, *Technological Change and the Growth of Health Care Spending* (Washington, DC: U.S. Congressional Budget Office, 2008).
- U.S. Food and Drug Administration, "Guidance for Industry: E9 Statistical Principles for Clinical Trials" (1998).
- "Novel New Drugs 2013 Summary" (2014).
- U.S. Government Accountability Office, "Medicare: High-Expenditure Part B Drugs" (2012).