Knowledge Growth and Specialization: Evidence from Oncologists *

Maya Lozinski
University of Chicago

November 2024

Abstract

Do expert workers respond to growth in general knowledge by specializing and narrowing their scope? I study this question empirically in the context of medical oncology. I document explosive growth in knowledge using changes in historical cancer treatment guidelines. Surprisingly, most oncologists are general oncologists and do not become more specialized. However, exposure to knowledge growth causes a minority of already specialized oncologists to become even more narrowly specialized. Over time, specialized oncologists start to prescribe more up-to-date cancer drugs and are more experienced with the drugs they prescribe than generalists suggesting they increasingly provide higher quality care than general oncologists. These results also suggest that most oncologists may respond to knowledge growth not by specializing but by falling behind the frontier. In addition, specialization only increases in large markets, implying growing economies of scale in specialized healthcare and leading to increased geographic inequality in local access to specialized care.

Keywords: specialization, knowledge growth, healthcare, economies of scale, technology adoption
JEL codes: I11, J24, L23, O33

*Thank you to Zarek Brot-Goldberg, Joshua Gottlieb, David Meltzer, Dan Black, Tamara Konetzka, Pauline Mourot, and Stephen Pardy. Thanks to the NIH Medical Scientist Training Program (T32GM007281), the NIA Program in Medicine, the Social Sciences, and Aging (T32AG051146), and the Center for Research Informatics (2U54TR002389-06). Author's email address: mayalozinski@uchicago.edu.

1 Introduction

Advances in fields such as computer science, biology, and medicine have lead to a rapidly growing knowledge base. Staying up-to-date with the frontier of the knowledge base is critical for many workers in these fields. However, the rapid rate of change means that staying up-to-date may be costly. Specialization, or narrowing one's scope, is one possible worker response to this growth, as this can reduce the cost of maintaining expertise. Workers may also face constraints to specialization, such as market size, coordination costs, financial disincentives, and other forces. Workers can either bear the growing cost of staying at the frontier across multiple areas or risk falling behind.

I evaluate how workers respond to growing general knowledge and assess the implications in the context of medical oncology. Medical oncologists are physicians who treat cancer patients with drugs (as opposed to surgery or radiation) and coordinate care across a team of cancer experts. They are individuals with extremely high human capital: they must complete college, four years of medical school, three years of clinical training in general internal medicine, and at least two more years of training in oncology to practice in the United States [1]. They are knowledge workers whose work largely involves cognitive and communication tasks. They are also primarily users of knowledge rather than producers of knowledge (Taylor, 2004; Jain et al., 2019; Schlafly and Sebro, 2022).

Despite large upfront investments in human capital, oncologists must continue to invest in human capital to stay at the frontier. Cancer is a condition with a vast demand for treatment — the lifetime risk of developing cancer is 40% (American Cancer Society, 2024)[2] —, leading to rapid innovation and knowledge growth. Per one news article, "incorporating this torrential evolution of knowledge is an impossible, Sisyphean task" (Wallis, 2023). As one oncologist told this author, "Some fields of cancer treatment are as complex now as the

[1] Note that other types of cancer specialists train in substantially different residency tracks than medical oncologists. For example, surgical oncologists first train in surgery for a particular region or structure in the body (e.g. the abdomen, for general surgeons) and then in cancer surgery for that region or structure.

[2] Fortunately, treatment is often life-extending or curative; the lifetime risk of dying of cancer is much lower at 20%.

whole of oncology was 20 years ago." As such, this setting enables us to study the effects of rapid knowledge growth in a context with material welfare implications.

Oncologists face a trade-off between the costs and benefits of specialization. For example, oncologists specializing in breast cancer may have more experience with breast cancer therapies than general oncologists, leading to higher quality of care or patient throughput. However, they may also face constraints that prevent them from specializing. For example, they may have difficulty finding enough breast cancer patients to fill their schedule. Patient volume may be particularly constraining in small markets, leading to greater specialization in large markets. Other types of constraints may bind regardless of market size. For example, oncologists give up income by specializing if breast cancer patients are less profitable to treat than other types of patients. Growing complexity in breast cancer treatment may shift oncologists towards becoming more specialized, but only in large markets, leading to growing geographic inequality.

First, I develop a novel measure of knowledge growth. I create a panel dataset with features of historical cancer guidelines using guidelines from the National Comprehensive Cancer Network (NCCN). The NCCN Clinical Practice Guidelines in Oncology (NCCN Guidelines ®) are the leading comprehensive cancer treatment guidelines in the United States. I measure the page length of NCCN Guidelines® for each type of cancer in each year. I then proxy growth in the clinically relevant knowledge base by cancer type as the increase in the length of associated guidelines over time. As a robustness check, I also measure changes in the number of citations to research papers and other sources in guidelines over time. A key advantage of guideline-based measures is that they capture the volume of knowledge directly relevant to the everyday practice of oncologists, as opposed to complexity or historical depth of the underlying science.

I confirm that oncology knowledge has grown dramatically. The total length of NCCN Guidelines, a "complete library" of cancer guidelines (NCCN, 2023*b*), has grown from 1,129 pages in 2002 to 5,825 pages in 2020. This growth in guideline length appears closely tied to

the publication of new scientific literature about treating cancer. It also occurs in a period with large increases in the number of available therapies: the total number of approved cancer drugs rose from 90 in 2002 to 243 in 2020.

Next, I measure oncologist specialization over time using a 21-year panel of Medicare claims data. Specialization is measured using an oncologist-level Herfindahl–Hirschman index (HHI) over time spent managing different types of cancer in a given year. The measure allows me to track subtle changes in specialization within an oncologist over time. I find that specialized oncologists have become even more specialized. Surprisingly, oncologists in general have not become more specialized. Median specialization declines slightly despite substantial growth in medical knowledge over this period.

Next, I estimate how much knowledge growth different oncologists are exposed to. I estimate an oncologist's exposure to knowledge growth as the increase in guideline length across all fields of cancer that they work in at baseline. This estimate of exposure to knowledge growth combines two sources of variation: variation in the sets of cancer subfields that physicians work in and variation in the rates of knowledge growth across those subfields.

One challenge is that this estimate of exposure to knowledge growth may suffer from omitted variables bias. For example, oncologists who work in more fields may be mechanically exposed to more knowledge growth and also differ systematically from oncologists who work in one or two fields. To address potential bias, I recenter the measure using the general approach developed by Borusyak and Hull (2020). The recentering removes variation in exposure to knowledge growth from physician scope (how many types of cancer a physician treats at baseline) and from predictable knowledge growth (predicted by how common a cancer is and whether it has surrogate endpoints). The remaining variation is attributed to "unpredictable" knowledge shocks, e.g., randomness and uncertainty in which clinical trials succeed or fail. Reassuringly, recentering removes the correlation between exposure to knowledge growth and observed omitted variables.

Focusing on 2008 to 2014, ex-ante specialized oncologists become significantly more spe-

cialized in response to knowledge growth. For every 100-page increase in relevant guidelines, these oncologists increase their specialization by approximately 200 HHI points, a 4% increase in HHI for a 13% increase in relevant guidelines. In contrast, ex-ante general oncologists do not become substantially more specialized in response to knowledge growth. Across most specifications, I can reject an increase in HHI of greater than 140 points in response to a 100-page increase in relevant guidelines or a 6% increase in HHI for a 10% increase in relevant guidelines.

Specialization increasingly appears to influence physician prescribing. Over time, specialists have become more experienced with the drugs they prescribe than generalists. They also prescribe newer drugs than generalists. These differences in drug prescribing have grown over time, consistent with growing quality differences between specialist and generalist oncologists.

The observed changes contribute to geographic inequality in specialization. The median oncologist in small and large markets is similarly general. In contrast, the 90th percentile of specialization within a market rises rapidly with market size, and this gradient has doubled over time, leading to growing geographic inequality in top specialization. As a result, local geographic access to highly specialized cancer has grown in the largest markets and stagnated or declined in small ones. These patterns also suggest that market-level economies of scale have grown over time for highly specialized medical care and labor markets in oncology.

My work contributes to several discussions in economics and healthcare. First, I contribute a novel measure of knowledge growth that comprehensively reflects currently relevant knowledge. Glied and Lleras-Muney (2008) measures technological change in healthcare using stocks and flows of drug approvals. The novel guideline-based measure overcomes some key limitations of using drug approvals. Specifically, guidelines include the full diversity of innovation and knowledge growth, which in oncology encompasses not only new drug approvals but also off-label use of drugs, new drug combinations, biomarkers, drug delivery strategies, and more. In addition, guidelines exclude information that is no longer pertinent, e.g., about

old drugs that are no longer used. An extensive literature in economics and business uses data on patents to study innovation and knowledge growth (Alcacer and Gittelman, 2006; Roach and Cohen, 2013). A key advantage of clinical guidelines over patent-based measures in this setting is that they enable the measurement of the volume of knowledge currently relevant to doctors' day-to-day work. For example, the chemotherapy drug sirolimus has a fascinating history that includes many patents and papers and is long enough to fill a book (Duffin, 2019). However, this history is largely irrelevant to the day-to-day clinical work of oncologists and not included in guideline-based measures of relevant knowledge.

Notably, I address a similar question to Jones (2009) about how knowledge growth impacts specialization but find differing results. Jones (2009) documents empirically that inventors face a growing historical "depth of knowledge" (as measured by patent citations) and that this leads them to specialize more. Similarly, in this setting, some oncologists become more specialized in response to knowledge growth. However, most oncologists do not. Users may bear the burden of knowledge more lightly than creators of knowledge like inventors as the use of knowledge may be less difficult than the creation of it. In addition, as I document, users of knowledge can fall behind the frontier when keeping up becomes costlier, an option that may not be available to creators.

The work contributes to an extensive empirical literature on the consequences of technological change in healthcare. Much of this literature focuses on single technologies with varied economic impacts on providers and patients (Skinner, 2011; Chandra, Malenka, and Skinner, 2014; Skinner and Staiger, 2015; Arrow, Bilir, and Sorensen, 2020). This work differs in considering the effect of changes in the overall corpus of technology. Evidence suggests that cumulative technological change in healthcare (as measured by the stock of drug approvals) increases educational inequality in health outcomes (Glied and Lleras-Muney, 2008). I build on this by examining how aggregate knowledge growth in a field impacts healthcare workers and geographic inequality in their work organization.

I also build on a small and emerging literature on economies of scale and agglomeration in

healthcare. Dingel et al. (2023) document that healthcare has large market-level economies of scale. I build on this by showing that economies of scale in healthcare appear to have grown over time and provide evidence that knowledge growth is a key mechanism.

2 Data and Measurement

The empirical strategy combines multiple data sources to measure knowledge growth and specialization in medical oncology over time. I measure changes in the page length and citations in cancer treatment guidelines over time and use this as a novel proxy for knowledge growth. I then use Medicare claims data to get physician-level estimates of exposure to knowledge growth over a five-year panel. I also use Medicare claims data to measure how physician specialization evolves over time. I supplement the analysis with additional information about physicians from the Medicare Data on Provider Practice and Specialty.

2.1 Cancer guidelines data

To measure growth in knowledge, I assemble a new dataset about past cancer guidelines from the National Comprehensive Cancer Network (NCCN). The NCCN is a prominent non-profit group composed of leading cancer centers (NCCN, 2023*a*). In 1996, they started to release treatment guidelines for common types of cancer and in 2001, they released the first "complete library" of cancer treatment guidelines (NCCN, 2023*b*)[3]. The NCCN Guidelines represent an expert consensus on approaches to cancer treatment. They are updated at least once a year (and often multiple times a year) to reflect the newest evidence and drug approvals. NCCN Guidelines are organized anatomically, with separate documents for different cancer types. They are widely used by oncologists; per one survey, 96% of oncologists think it is important to use NCCN Guidelines in patient care (McGivney, 2008).

The NCCN Guidelines contain recommendations for all stages of the cancer treatment

[3]Note that a few rare cancers lack guidelines after 2001. See Appendix Section B.1 for further discussion.

process, including diagnosis, drug treatments, radiation, and surgery (McGivney, 2008). Medical oncologists are trained in comprehensive care cancer care management; they are responsible both for managing "systemic therapies" such as chemotherapy, and serving as the central node in the network of other specialists who may be involved in a patient's cancer care (Popescu et al., 2014).

The second role is to serve as the central node in a network of other specialists who may be involved in a patient's cancer care, e.g., surgical oncologists, radiation oncologists, and pathologists.

Changes in NCCN Guidelines are used as a proxy for changes in volume of relevant clinical knowledge. I measure the change in the page length of the first set of NCCN Guidelines for each year for each cancer type, and use this to proxy the change in the volume of relevant knowledge by cancer type. As a robustness check, I also use the change in the number of references in NCCN Guidelines.

Appendix Figures A.1 and A.2 includes examples of guideline pages from 2009 to 2014. Appendix Figures A.1a and A.1b show the table of contents for non-Hodgkins lymphoma (NHL). Between 2009 and 2014, the length of the table of contents grows longer and includes new subtypes of NHL. Appendix Figures A.2a and A.2b show a flowchart with the diagnostic workup for lobular carcinoma in-situ, a subtype of breast cancer. Between 2009 and 2014, the flowchart grows more complex and links and links to five additional flowcharts, up from three.

To further characterize the nature of knowledge growth, I also further examine changes in the references in NCCN Guidelines. I digitize the references text in the NCCN Guidelines for 2009 and 2014 and use this text to link references to published scholarly works in OpenAlex, a database of research publications (Priem, Piwowar, and Orr, 2022).

Several key assumptions are required for change in page length to be a valid proxy for the change in the volume of knowledge. First, it assumes that each page represents an equal volume of knowledge. Reassuringly, the formatting of the NCCN Guidelines appears

was unchanged and standardized across cancer types over the main time period for the research design (2008-2014), suggesting that longer page lengths result from more content, not changes in formatting (e.g. increases in font size). In fact, changes in page lengths are likely to slightly understate, not overstate, knowledge growth. Appendix Figures A.1 and A.2 both illustrate cases where each page contains more content over time.

The measure also assumes that any relevant knowledge not included in the NCCN Guidelines is either exactly proportional to the volume of information the guideline pages contain or is constant across cancer types. For example, the NCCN Guidelines may not include all relevant information about a chemotherapy drug, but they need to contain a consistent share of the information about each drug.

Lastly, growth in page count also assumes knowledge growth is primarily additive. For example, if a 20-page guideline grows to 30 pages, I assume 10 pages of new knowledge were added. If, in fact, 20 pages of new knowledge have been added, and 10 pages of old knowledge have been dropped, I will understate the amount of knowledge growth. I assess this assumption empirically in section 3.

This proxy measure provides several key advantages over other measures of knowledge growth. For example, alternative approaches include using new journal publications, patents, or drug approvals by cancer type. However, many of these new papers or patents may be irrelevant to the day-to-day practice of medicine. In addition, older papers or drugs may remain relevant. For example, many (but not all) citations from the 2002 NCCN guidelines are also included in the 2020, suggesting some medical evidence has a long shelf life.

2.2 Medicare claims data

I use a 21-year panel of Medicare claims data to measure oncologists' work patterns. Specifically, I use a 20% sample of Medicare Part B Claims from 1999-2019. Medicare provides federally sponsored health insurance coverage for disabled adults and most U.S. adults over 65. This age group bears the highest cancer incidence in the U.S., making it the ideal

population to study cancer care (for Disease Control and Prevention, 2019).

I construct an annual panel of oncologists. I restrict the sample to physicians whose most commonly listed specialty in a year is Hematology-Oncology, Hematology [4], or Medical Oncology. I further restrict the sample to physicians who evaluated and managed at least 25 unique cancer patients in the data in a year. This restriction ensures that the measures of specialization are based on a minimum patient volume. I consider how patient volume impacts the measurement of specialization and the overall results extensively in Section 5.3 and Appendix C. For the main regression analysis, I subset to a balanced panel of oncologists from 2008 to 2014.

Next, I determine which types of cancer an oncologist manages. For this, I focus on the outpatient evaluation and management visits billed by each oncologist, which limits the sample to instances where the oncologist managed long-term care and decision-making for a patient, which requires high-level medical knowledge and judgment (as opposed to, for example, administering a chemotherapy that someone else prescribed or managing acute medical issues for hospitalized cancer patients)[5]. If an oncologist sees a breast cancer patient, they are denoted as having managed breast cancer[6].

I approximate the time each oncologist spends managing each type of cancer, which I later use to measure physician specialization. I map each evaluation and management code to work Relative Value Units (wRVUs), a measure used by Medicare to estimate the amount of physician work the code requires. For evaluation and management codes, wRVUs reflect the expected time required.

[4] Hematology includes the treatment of blood cancers.

[5] As a robustness check, I alternatively restrict the sample to new patient visits in section **??**.

[6] A patient's cancer type is defined in the analysis as the specific type of cancer for which they have the most diagnosis codes in a year. For example, if a patient has a diagnosis code for cervical cancer billed once and diagnosis codes for breast cancer billed ten times in 2008, they are considered to be a breast cancer patient in 2008. If a patient has diagnosis codes for both "Unspecified" cancer and breast cancer, they are also considered a breast cancer patient (regardless of how many "unspecified" codes are billed). Patients with only unspecified or non-specific cancer types are excluded from the sample (1.9% of patients in 2009).

2.3 Medicare Data on Provider Practice and Specialty

I also use the Medicare Data on Provider Practice and Specialty (MD-PPAS) to observe physician characteristics, namely their birth year and the top two tax ID numbers from which they bill. The birth year allows me to determine physician's age at the start of the panel period, which is useful because age may influence specialization. MD-PPAS data also contains the top two Tax ID numbers from which each oncologist bills. Multiple physicians can share the same tax ID number, which implies they are in the same organization or group practice. I use shared tax ID numbers to estimate the number of oncologists at an organization because organization size may also influence specialization[7].

2.4 Measuring physician specialization

The main outcome of interest is oncologist specialization. Specialization is measured using a Herfindahl–Hirschman Index (HHI) over the shares of time that an oncologist spends managing each type of cancer. Denote T_{ict} as the time (proxied by wRVUs billed in Medicare) that oncologist i spends managing cancer $c \in C$ in year t. HHI is calculated as follows:

$$HHI_{it} = 10,000 \sum_{C} \left(\frac{T_{ic}}{\sum_C T_{ic}} \right)^2$$

In this definition, HHI_{it} has a maximum value of 10,000, which is achieved if an oncologist spends all their time in year t managing one type of cancer. If an oncologist divides their time equally between two types of cancer, then HHI_{it} is 5,000. The lower limit of HHI_{it} is 0, which is approached as an oncologist spends an infinitesimally small amount of time on an infinite number of cancer types.

This measure has several key advantages. First, HHI is continuous and based on the work that an oncologist performs in a year. Second, it is sensitive to subtle intensive margin

[7] A limitation of this measure is that healthcare organizations may have multiple tax ID numbers or providers within the organization may all bill to personal tax ID numbers. If so, the number of physicians at the tax ID may understate how many people are at an organization. As such, this variable provides useful but imperfect information about if a physician is a solo provider or in a bigger organization.

changes in specialization. The measure will capture if an oncologist shifts towards focusing on one particular type of cancer, even if they do not entirely stop seeing long-standing patients for other types of cancers. As such, it can capture differences in specialization for the same worker over time and across workers with identical job titles and specialty fields.

2.5 Measuring relevant knowledge growth for each oncologist

To estimate how knowledge growth impacts specialization, I combine the Medicare claims and NCCN Guidelines to construct a measure of the growth in relevant knowledge for each oncologist. For each oncologist i, I denote the set of cancer types they manage based on 2008 Medicare claims data as C_i. Each cancer, c, has an increase in relevant guideline page count Δk_c over the following five years (2009 to 2014). I combine these values to estimate the knowledge growth in the fields that oncologist i worked in: $\Delta k_i = \sum_{C_i} \Delta k_c$. As a robustness check in section **??**, I raise the threshold for a physician to be considered working in a field by ignoring cancers where they manage only one or two patients a year in the data. To illustrate, suppose oncologist i treats only lung and breast cancer in 2008 and that over the following five years, the lung cancer guidelines increased by 50 pages ($\Delta k_{lung} = 50$) and the breast cancer guidelines increased by 100 pages ($\Delta k_{breast} = 100$). Then there are 150 pages of knowledge growth over the following five years (2009 to 2014) in the fields that oncologist i worked in.

To link between cancers treated and guideline page lengths, I make a set of custom diagnosis code and cancer guideline groupings. Because there are discrepancies between NCCN Guidelines and standard diagnosis code groupings, I create custom groupings based on the anatomical origin of the cancer to ensure stable categories over time[8]. For example, the NCCN Guidelines for basal and squamous cell cancer (both types of skin cancer) were

[8]Categorization based on anatomical origin mirrors the organization structure of NCCN Guidelines, specialization, and cancer treatments in this time period. In recent years, some drug approvals have been given based on biomarkers without regard to anatomical origin of the cancer. However, the very first such approval was given in 2017 (Olivier, Haslam, and Prasad, 2021), outside the key panel period studied.

contained in a single guideline document until 2013 and then divided into two separate documents starting in 2014. I group these and other related NCCN Guidelines together into the overarching category of skin cancer guidelines. More detail about these groupings can be found in appendix section B.2.

2.6 Physician prescribing data

To understand how physician prescribing evolves, I measure various features of the typical bundle of anti-cancer drugs that an oncologist prescribes in a year. Anti-cancer drugs are defined using World Health Organization Anatomical Therapeutic Class (ATC) codes that fall into the list of anti-cancer drugs from Pantziarka et al. (2021). I include only one code per drug, even if a drug has multiple ATC codes, representing multiple medical uses. Drug age is calculated as years since the first FDA approval for the drug for any indication[9]. I use data on oncologists' prescriptions from a 20% sample of Medicare Part B and Part D claims from 2008 to 2019[10]

The sample changes slightly from the main panel analysis. In addition to the 25 patient minimum, the sample is restricted to oncologists who prescribed anti-cancer drugs to at least 11 unique patients across a 20% sample of Part B and Part D data. This restriction ensures drug bundle features are calculated over a sufficient volume of patients per oncologists to meet privacy requirements and permit more meaningful inference. To counteract this reduction in sample size, I examine the full, rather than balanced, panel of oncologists when describing prescribing patterns. The use of the full panel enables examination of prescribing patterns over a longer time horizon (2008-2019) with a larger number of oncologists than using the balanced panel.

[9]This approach raises the average drug age relative to a strategy that accounts for new indications. However, I cannot observe which indication a drug was given for or if it was given off-label, without an FDA-approved indication, and so use the year of the first indication.

[10]2008 is the first year in which both Medicare Part B and D claims use National Provider Identifiers to identify the majority of prescribers, enabling linkage between the two parts.

3 Descriptive Patterns

This section describes the characteristics of oncologists in the sample and the patterns of knowledge growth and specialization over time.

3.1 Knowledge growth over time

The volume of clinical knowledge, as proxied by guideline page lengths, has grown substantially over time. This growth in page length is driven by growth in the relevant scientific literature, which is largely additive. In addition, the referenced scientific literature is almost entirely distinct across guidelines, suggesting the underlying knowledge is also distinct.

Cancer treatment guidelines have increased over five-fold in length from 2002 to 2020, as shown in Figure 1a and Appendix Table A.1. One factor likely contributing to this growth is the explosion in the number of drugs available to treat cancer, as shown in Figure 1b[11]. Not only has the total number of drugs been increasing rapidly over time, but the rate at which new drugs are introduced has been accelerating.

An additional factor in guideline growth is the addition of new scientific literature to guidelines. Figure 1c shows that the number of references in these guidelines has also multiplied over time. This growth in references appears to drive the growth in page length; Table 1 shows that guideline page length is tightly correlated to the number of references both in cross-section and over time. Over time, every additional reference is associated with an additional $\frac{1}{5}$th of a page of guidelines, of which $\sim \frac{1}{15}$th of a page is the direct text of the reference[12] and $\sim \frac{2}{15}$th of a page is additional guideline contents.

The vast majority of references are to published research papers. Of 14,389 references in 2009 and 2014, 78% can be linked with high confidence to published scholarly works. The remaining 22% include drug package labels, conference posters, professional society issue briefs, government reports, and papers that cannot be linked with high confidence.

[11] Data on anti-cancer drugs is from Pantziarka et al. (2021) and uses Anatomical Therapeutic Chemical Level 5 classifications to define individual drugs.

[12] Guidelines from this period fit about 15 references on a page

Knowledge appears to overlap minimally across cancers. In 2009 and 2014, 94% of cited scholarly works are cited by only one set of guidelines.

Knowledge growth appears to be largely additive, though there is modest degree of turnover. The 2009 guidelines reference 3,177 unique scholarly works, of which 2,071 continue to be referenced in 2014. An additional 5,195 unique scholarly works are added by 2014, greatly outpacing the 1,106 which are dropped. Figure 2 shows that these additions are largely new or recent scholarly works.

3.2 Oncologist characteristics

This section describes summary statistics for the balanced panel of physicians used in the main analysis. Table 2 shows summary statistics for the 3,948 oncologists with sufficient patient volume in the data (25 patients) in the key panel years of 2008, 2009, and 2014. On average, they see 61 unique cancer patients in the outpatient setting in the data in 2008. The number of distinct cancer types each oncologist treats ranges from 1 to 25, with an average of 14. The mean oncologist HHI in 2009 is 2,920, with a maximum of 10,000, demonstrating that some oncologists treat only one type of cancer. Their average age at the beginning of the panel is 49 years old.

In 2009, the total length of NCCN Guidelines was 1,832 pages, though only 1,086 pages are relevant to each oncologist, on average. Over the following five years, the relevant guidelines per oncologist (holding constant the set of relevant cancers) nearly double, to 2,118 pages.

3.3 Trends in oncologist specialization over time

To understand how specialization is evolving, I characterize the changing distribution of oncologist specialization over time. Figure 3 shows the 10th, 25th, 50th, 75th, and 90th percentile of oncologist HHI between 1999 and 2019. It also plots the HHI of the field overall, which is calculated using the distribution of cancer types seen by oncologists in

aggregate. Several striking trends stand out.

First, there is a sharp, discontinuous decline in HHI at all points in the distribution between 2014 and 2016, corresponding to the switch from ICD-9 to ICD-10 diagnosis coding midway through 2015. This clear trend break motivates the decision to end the main balanced panel in 2014, which prevents coding changes from contaminating the estimates of change in HHI.

Second, the typical oncologist did not become substantially more specialized over time. Median specialization and below is flat or declining slightly over this period. The distribution's 10th, 25th, and 50th percentiles are all close to each other and the HHI of the field in aggregate. This result suggests that most oncologists are "general" oncologists who see a similar distribution of cancer patients as the field overall. Given the several-fold growth in knowledge and treatment options for cancer over these 21 years, the lack of growth in specialization for the median oncologist is surprising.

In contrast, subspecialized oncologists have become dramatically more specialized over time. The 90th percentile of oncologist HHI increases by 1,030 HHI points from 1999 to 2014 and another 298 points from 2016 to 2019.

As a robustness check, Figure 3 is replicated using 25 random patients per oncologist, rather than all patients, to measure HHI in Appendix Figure A.3. This removes any variation in patient volume that might contribute to the observed trends. The magnitude of the changes is attenuated, but the trends are the same, with the 90th percentile of specialization rising and the 10th-50th percentiles of specialization declining.

These time trends suggest the typical oncologist keeps a broad clinical practice even in the face of rapid advances in cancer care. However, the rise in extreme specialization leaves open the possibility that some specialized oncologists respond to these rapid advances by narrowing their clinical focus even further in order to stay at the frontier. To tease out these differences, the remaining analysis considers patterns of specialization and technology adoption separately for more general and more specialized oncologists.

3.4 Differences between more and less specialized oncologists

In the panel data, specialists and generalists have different baseline characteristics, as shown in Table 3. Generalists are defined as those below the 75th percentile of physician HHI in the first year of the panel, whereas specialists are those with an HHI at or above the 75th percentile. In 2008, specialists have an average HHI of 4,866, double that of the generalists, who have an HHI of 2,269 on average. Specialists are far more likely to be academic physicians (see Appendix B.3 for more details on how I proxy academic status using Tax IDs). Specialists are also more concentrated in large markets: 60% of them are in the most populous 20% of hospital referral regions (HRR) as compared to 47% of general oncologists. In addition, they work in larger organizations: On average, specialists work in places with 21 oncologists, compared to 17 for generalists. These empirical differences are consistent with large firms, large markets, and academic institutions all facilitating specialization.

Notably, age does not differ meaningfully between specialists and generalists (50 versus 49 years old at the beginning of the panel). This lack of age difference is consistent with how specialization often works in medical oncology: many oncologists decide whether to become generalists or specialists while in their fellowship training.

4 Identification

To test whether knowledge growth leads to specialization, physician-level changes in specialization are regressed on measures of exposure to knowledge growth. To account for non-random exposure to knowledge growth, these measures are "recentered" to remove variation from highly predictable knowledge growth using the general approach from Borusyak and Hull (2020).

4.1 Estimating Equation

The main estimating equation is as follows. Denote physician i's change in specialization, as measured by HHI, between 2009 and 2014 as ΔHHI_i. Denote their exposure to knowledge growth as Δk_i. Recall that exposure to knowledge growth is calculated as $\Delta k_i = \sum_{C_i} \Delta k_c$, where C_i is the set of cancers that physician i treats in 2008 and Δk_c is the change in page length for each type of cancer between 2009 and 2014. Also, denote the vector of other physician characteristics as X_i. These include physician age, market, patient volume, and organization size at the beginning of the panel. The main estimating equation is as follows:

$$\Delta HHI_i = \alpha + \beta \Delta k_i + \theta X_i + \epsilon_i$$

This equation is estimated using ordinary least squares with robust standard errors. I observe individual physicians over time, so the design captures within-physician physician variation in specialization. The design also allows for specialization to respond immediately to knowledge growth. If the specialization response is highly lagged, that may attenuate the estimates.

The main coefficient of interest is β, the effect of knowledge growth on specialization. β is estimated using variation in exposure to knowledge growth (Δk_i), which comes from two sources. The first source is that physicians work in different sets of cancer subfields, e.g. some treat breast cancer, some treat leukemia, and some treat both. The second source of variation is that different cancer subfields have different amounts of knowledge growth. For example, leukemia may have more new drugs and therapeutic advances than breast cancer, so physicians who treat only leukemia are more exposed to knowledge growth than those who treat only breast cancer. The intercept term (α) absorbs any global shocks to knowledge that affect all fields of cancer. As such, the main coefficient of interest (β) captures the effect of changes in different cancer subfields, rather than global shocks.

4.2 Recentering instrument to correct for omitted variables bias

The measure of exposure to knowledge growth may suffer from omitted variables bias because physicians non-randomly choose which fields to work in. To correct for potential bias and estimate causal coefficients, I recenter the measure using the general approach from Borusyak and Hull (2020).

In this context, the omitted variables bias may arise because physicians choose which fields to work in and those who work in more fields will mechanically be exposed to more knowledge growth. For example, someone who treats both breast cancer and leukemia will mechanically be exposed to more knowledge growth than someone who treats only breast cancer. They may also be more likely to work in a small healthcare market that doesn't have the patient volume to support a breast cancer specialist. Small markets may also have different trends in specialization than large ones for reasons outside of knowledge growth e.g. more hospital consolidation. As a result, the change in specialization (ΔHHI_i) may be correlated with exposure to knowledge growth (Δk_i) because both are correlated with market size, leading to omitted variables bias in the estimate of β.

Empirically, realized exposure (Δk_i) is calculated from the data. Conceptually, it can also can be decomposed into two components.

$$\Delta k_i = \underbrace{\mathbb{E}[\Delta k_i]}_{\text{Expected Exposure}} + \underbrace{\Delta \tilde{k}_i}_{\text{Remaining Exposure}} \tag{1}$$

The first component, $\mathbb{E}[\Delta k_i]$, is the expected exposure to knowledge growth and contains the omitted variables bias. The second component, $\Delta \tilde{k}_i$, is the remaining exposure to knowledge growth. This component is exogenous and uncorrelated with other factors that impact specialization except for knowledge growth.

The realized exposure (Δk_i) can be substituted with the "recentered", or residual, expo-

sure ($\Delta\tilde{k}_i$) for an unbiased and efficient estimate of β (Borusyak and Hull, 2020), as below:

$$\Delta HHI_i = \alpha + \beta\Delta\tilde{k}_i + \theta X_i + \epsilon_i \tag{2}$$

The formal identification assumption here is that recentered exposure to knowledge growth is uncorrelated with the contents of the error term, including omitted variables ($\Delta\tilde{k}_i \perp \epsilon_i$). The real-world content of this mathematical assumption depends on exactly how one forms the expectation of exposure to knowledge growth.

I broadly formulate the expectation as follows. Assume there exists an expected knowledge growth for each type of cancer, $\mathbb{E}[\Delta k_c]$. Then a physician's expected exposure to knowledge growth ($\mathbb{E}[\Delta k_i]$) can be calculated as below:

$$\mathbb{E}[\Delta k_c] = \sum_{C_i} \mathbb{E}[\Delta k_i] \tag{3}$$

The central challenge is to formulate a plausible estimate of expected knowledge growth by cancer type, $\mathbb{E}[\Delta k_c]$.

The simplest assumption is that all cancer fields have no expected knowledge growth ($\mathbb{E}[\Delta k_c] = 0$). In this case, recentered exposure ($\Delta\tilde{k}_i$) equals realized exposure (Δk_i). This is the underlying assumption built into using the original measure of exposure to knowledge growth (Δk_i). The assumption that doctors do not anticipate any knowledge growth seems implausible to me and liable to lead to residual omitted variables bias as discussed above.

A somewhat more plausible assumption is that all cancer fields have the same non-zero level of expected knowledge growth. Doctors expect knowledge to grow in all fields, but do not anticipate which ones will have more growth than others. The assumption that all cancers have the same level of expected knowledge growth can be operationalized as $\mathbb{E}[\Delta k_c] = \overline{\Delta k_c}$ where $\overline{\Delta k_c}$ is the mean realized page increase across cancer types. This estimate of $\mathbb{E}[\Delta k_c]$ can then be used to back out the recentered measure of exposure to knowledge growth ($\Delta\tilde{k}_i$) using Equations 1 and 3. This approach compares doctors who treat the same number of

cancers and is similar to controlling for the number of cancers an oncologist treats (Borusyak and Hull, 2020).

Critically, this approach assumes that cancer-level knowledge shocks are unpredictable and uncorrelated with physician sorting, which may not be the case. For example, an oncologist may anticipate that treatments for breast cancer will advance rapidly and thus choose to specialize in breast cancer. Addressing this potential bias motivates further refinement to estimating expected knowledge growth.

Based on findings in the literature, I allow expected knowledge growth ($\mathbb{E}[\Delta k_c]$) to vary with cancer characteristics. Acemoglu and Linn (2004) find that conditions with more potential patients had more innovation, suggesting more common cancers may have more knowledge growth. Budish, Roin, and Williams (2015) find that cancers have more innovation if they have high-quality surrogate endpoints for mortality in clinical trials, since they allow researchers to more quickly and easily determine if a cancer is responding to a treatment.

Let p_c denote the number of patients in Medicare with each type of cancer in 2008, a measure of the initial prevalence of the cancer. Also, define e_c as an indicator for if a cancer type has surrogate endpoints per Budish, Roin, and Williams (2015). I include both of these in a Poisson maximum likelihood regression to predict the expected page count increase for each type of cancer:

$$E[ln\Delta k_c] = \delta_0 + \delta_1 e_c + \delta_2 log(p_c) + \delta_3 e_c log(p_c) \tag{4}$$

Note that this equation also accounts for knowledge shocks shared across all cancer types (captured by δ_0). Table 4 shows the coefficient estimates for the above regression. As above, I use the resulting estimate of expected knowledge growth by cancer type ($E[\Delta k_c]$) to back out each physician's recentered exposure to knowledge growth ($\Delta\tilde{k}_i$), removing variation from both physician scope and cancer characteristics.

This approach entails several key assumptions. First, it assumes that the above factors

capture most of the predictable variation in knowledge growth across cancer types. Empirically, this assumption is supported by the observation that these factors alone explain 77% of the observed variation. I attribute the remaining variation largely to randomness and uncertainty in the research and development process, e.g., some clinical trials and lines of drug development fail while others succeed, and assume they are largely unrelated to oncologists' specialization decisions.

4.3 Correlation of Exposure to Knowledge Growth with Omitted Variables

I find that the original measure of exposure to knowledge growth strongly correlates with observed omitted variable such as market size suggesting recentering is necessary for unbiased estimation. Fortunately, recentering with progressively more sophisticated estimates of expected knowledge growth ($\mathbb{E}[\Delta k_c]$) reduces this correlation until it is largely eliminated.

Following the strategy of (Borusyak and Hull, 2020), I test in Table 5 if the measures of exposure to knowledge growth correlate with observed physician characteristics including age, patient volume, firm size, and market size. Column 1 reveals that these physician characteristics are strongly correlated with the unadjusted measure, explaining 33.5% of the variation. For example, an increase in market size of 1 million people is associated with a decrease in exposure to knowledge growth of -15 pages ($\sigma = 2$). These results substantiate concerns about omitted variables bias; a regression using this unadjusted measure will compare physicians in larger and smaller markets.

Column 2 removes physician scope variation, Column 3 removes cancer prevalence variation, and Column 4 removes surrogate endpoint variation, respectively. Addressing each source of omitted variables bias reduces the explanatory power of physician characteristics until all significant correlation has been removed. The observed physician characteristics explain only .1% of the total variation in the fully recentered measure as compared with 33.5% of the variation in the original measure, a decrease of over 99%. Appendix tables A.2

and A.3 show the same regressions separately for specialist and generalist oncologists, with substantively identical results.

5 Effect of Knowledge Growth on Specialization

Using the recentered measure, I find that initially specialized oncologists become more specialized in response to knowledge growth, but general oncologists largely do not.

5.1 Effects on specialization

Figure 4a shows that specialized oncologists become even more specialized in response to knowledge growth. For every 100-page increase in relevant guidelines, their HHI increases by 196 HHI points ($\sigma = 80$) in the univariate regression. Relative to baseline (2008) means, this represents a 4% increase in HHI for a 13% increase in relevant guidelines. Table 6 shows that the estimated effect increases slightly to 246 ($\sigma = 91$) after adding fixed effects for physician age, hospital referral region, patient volume, and firm size. These results are surprising because, at the outset, these specialists only need to keep up with a reduced subset of medical advances relative to generalists. Yet when faced with further medical advances, they narrow their clinical practice even more.

In contrast, Figure 4b reveals that general oncologists do not become substantially more specialized in response to knowledge growth. Their response to knowledge growth is statistically indistinguishable from zero, with an estimated increase in HHI of 51 points ($\sigma = 31$) for every 100-page increase in relevant guidelines in the univariate regression. Adding fixed effects in Table 7 increases the coefficient estimates in some specifications to just above the traditional threshold for significance. However, I interpret these results as most consistent with a null finding, given than any statistical significance depends on the fixed effects used. Across specifications, I can reject an increase in HHI of greater than 140 points in response to a 100-page increase in relevant guidelines, or a 6% increase in HHI for 10% increase in

relevant guidelines. The typical oncologist appears largely undeterred from continuing a broad general practice by rapid changes in medical treatments.

For specialists, the increase in specialization does not appear to occur on the extensive margin, suggesting it is driven by the intensive margin. Table 8 shows specialists do not significantly reduce the count of unique cancers that they treat in a year. In practical terms, specialized oncologists focus their practice more on some types of cancer but do not entirely stop seeing other types. This pattern may reflect specialist oncologists continuing to see established patients but limiting their new patients to those with a narrower set of cancers. A relative paucity of new patient visits makes it difficult to assess this hypothesis with statistical precision (Table 10).

Generalists may slightly reduce their scope of practice on the extensive margin. Table 9 shows that generalists drop about .2 unique cancers from their patient panel per 100-page increase in guidelines. This decrease is significant in most specifications but also sufficiently precise to rule out economically meaningful effects. Across specifications, the estimates rule out a reduction in the unique count of cancer types managed by generalists of .5 or more (on a baseline mean of 15). A few generalists may drop a cancer type from their practice but such a response is minimal.

Appendix Figure A.4, Table A.4, and Table A.5 show pooled estimates for the above regressions and find results consistent with averaging the effects for specialists and generalists. Collectively, these findings imply that growing medical knowledge leads oncologists to become slightly more specialized in aggregate, driven a growing right tail of extreme specialization.

5.2 Robustness to alternative measures of knowledge growth

The findings are robust to alternative assumptions about which patients are relevant to oncologists and to an alternative measure of knowledge growth using guidelines references.

One concern is that physicians may respond to technological change only in their main

areas of clinical practice. They may not change their scope of practice in response to advances for cancers that they see only occasionally. If the case, the main measure of knowledge growth may include a lot of irrelevant variation. Table 10 estimates again the effect of knowledge growth on specialization for specialists but ignores knowledge growth in cancer types where physicians saw only one or two patients a year in the data. The estimated effect remains quite similar ($\beta = 310, \sigma = 110$). Table 11 repeats this analysis for generalists and continues to estimate no response to knowledge growth ($\beta = 9, \sigma = 35$). This robustness may reflect the fact that the data includes only a 20% sample of Traditional Medicare patients. A physician who see one breast cancer patient in the data is expected to see five breast cancer patients in Traditional Medicare, and potentially many more in Medicare Advantage, Medicaid, or private insurance.

I also consider the robustness of the main regression results to an alternative measure of knowledge growth — the count of sources referenced in each guideline. This alternative measure reflects the quantity of academic literature and other sources informing the guidelines. Substantively, the results are similar. More cited sources leads specialists to become significantly more specialized (Appendix Table A.9). More cited sources causes smaller increases in generalists degree of specialization (Appendix Table A.10) In aggregate, more cited sources has an intermediate and positive estimated effect on specialization (Appendix Table A.11)

5.3 Robustness to volume-dependent measurement error

Lastly, I consider how measurement error in the measure of specialization ($\hat{HHI}_i$) impacts the main findings. Appendix C addresses this issue in greater detail. Here, I summarize the key points.

The data does not include all patients seen by an oncologist which may lead to measurement error in both the outcome variable of specialization ($\hat{HHI}_i$) and the treatment variable, recentered exposure to knowledge growth ($\Delta\tilde{k}_i$). In particular, these quantities may be measured with systematically more measurement error for physicians with lower pa-

tient volumes. Reassuringly, I find that my main treatment variable, recentered exposure to knowledge growth ($\Delta\tilde{k}_i$), is uncorrelated with patient volumes, implying it is uncorrelated with volume-dependent measurement error. My main outcome variable, specialization, is measured with patient volume-dependent measurement error. However, the main regression results are essentially unchanged after adding fixed effects for patient volume (Tables 6 and 7), assuaging concerns of bias.

As an additional set of robustness checks, I replicate the key descriptive and regression results with 25 random patients per physician instead of all observed patients. This approach inflates measurement error for all physicians who see over 25 patients, more so for physicians with many patients, equalizing measurement error across physicians. With this change, the time trends in specialization (Appendix Figure A.3) are somewhat attenuated but show a similar rise in top specialization with declining 50th percentile and below specialization. With the additional measurement error, the estimates of the effect of knowledge growth on specialization are noisier but generally consistent with more specialization in response to knowledge growth (Appendix Tables A.6, A.7, and A.8).

6 Technology Use and Specialization

In this section, I focus on the anti-cancer drugs that oncologists use. As would be expected, oncologists are the most knowledgeable physicians about the use of these anti-cancer drugs. These drugs are complex: they often have complicated dosing regimens and serious toxicity risks. Knowing how to treat patients with these risky drugs is key to cancer treatment, and given the morbidity and mortality risks, examining oncologists' prescription patterns can provide important insight into the quality of patient care. I find that specialized oncologists use fewer unique drugs and have more experience in each drug they prescribe. They also tend to prescribe newer drugs to patients. Importantly, these differences between specialist and generalist oncologists have grown substantially over time.

Figure 5 shows that from 2008 to 2019, specialists and generalists increasingly diverged in the terms of the number of unique drugs they prescribed in a year. Over this period, the most general oncologists (1st—50th percentile of HHI) prescribe a similar number of unique drugs each year. In contrast, the most specialized oncologists (75th - 99th percentile of HHI) steadily reduce their scope. This leads to growing divergence in the scope of technology use between these groups. In 2008, specialists prescribe 9% fewer unique drugs than generalists, growing to 27% fewer drugs in 2019. Appendix Figure A.5 finds that these patterns persist after removing differences in patient volume by degree of specialization. These findings indicate that specializing allows oncologists to narrow the scope of technologies that they use.

Using fewer unique drugs enables specialists to have more experience with the drugs they do prescribe. Figure 6 shows oncologists' relative experience with the drugs they use, with experience measured as the count of patients prescribed to in a year. I focus on experience with a drug within a year since physician human capital appears to decay rapidly (Hockenberry and Helmchen, 2014; Facchini, 2022). In 2008 the most and least specialized oncologists have similar levels of experience with the drugs they prescribe. This steadily diverges over time as specialists become increasingly more experienced than generalists. By 2019, specialized oncologists are approximately 20% more experienced with the drugs they use than general oncologists. As these drugs are typically complex and high-risk, more experience with a drug may translate into better clinical management of drug dosing, toxicities, and complications from use, leading to better patient outcomes.

Specialization also appears to benefit oncologists by enabling them to better keep up with advances in cancer therapies. Figure 7 displays average age of oncologists' drug bundles over time, with each drug weighted by how many patients the oncologist prescribed it to. In 2008, specialist and generalist oncologists prescribe drugs of a similar age. By 2019, specialist oncologists prescribe drugs that are significantly newer — about two years newer on average. These results are consistent with specialization enabling oncologists to stay at the frontier

of medical advances and adopt new drugs more quickly.

In fact, generalist oncologists initially appear not to update their drug bundle. Their drug bundle ages one year every year. This result is driven by the use of steroids, an old medication class, which are given to many patients and hence weighted heavily in the drug bundle. Appendix figure A.7 drops steroids. As before, specialists start to update their drug bundle faster over time than generalists, and pull ahead in the use of new technology. Generalist oncologists do update their prescribing, but not fast enough to keep up with the frontier.

Collectively, these results show that specialists have increasing different prescribing patterns from generalists. They also imply that specialized oncologists may provide higher quality care than generalists and that these quality differences may be growing over time. Unfortunately, the data does not contain sufficient detail about each patient's cancer (e.g., stage, biomarkers) to directly test if specialists provide more guideline concordant care or improve patient outcomes. However, several case studies from the medical literature find that academic centers provide more guideline concordant care and are associated with improved patient survival (Voss et al., 2017; Thiels et al., 2020). As specialists are concentrated in academic centers, these results are consistent with the hypothesis that specialists provide higher-quality care and improve patient survival.

7 Constraints to Specialization

Why do specialist and generalist oncologists respond differently to knowledge growth in terms of specialization? To unpack this, I assess the constraints to specialization for specialists and generalists in the context of classic theories of specialization from Smith (1776) and Becker and Murphy (1992). General oncologists appear to constrained by some combination of low returns to specialization and high coordination costs. In contrast, specialist oncologists encounter market size constraints.

7.1 Theories of constraints to specialization

I focus on two key potential constraints to specialization: market size and coordination costs. Smith (1776) was the first in a long line of many economists to observe that specialization is often limited by the extent of the market. Under this constraint, specialization grows with market size, as shown in Figure (a). Becker and Murphy (1992) highlight a contrasting mechanism and prediction: coordination costs between specialized workers may limit specialization independent of market size. Specialization may make individual workers more productive at specific tasks. However, this splits tasks among more workers, leading to higher coordination costs, e.g., communication overhead or incentive compatibility challenges. If coordination costs are a binding constraint, then specialization will remain flat across market size, as also shown in Figure (a).

I investigate which of these predictions is most consistent with the data for both specialists and generalists. I measure market size as the population of the oncologist's main hospital referral region and consider both the median and right tail of specialization within a market.

Figure 8c and Table 12 show that top specialization, defined as the 90th percentile of physician HHI within a market, rises substantially with market size. In 1999, doubling market size leads to a 495 point increase ($\sigma = 79$) in the 90th percentile of HHI. This slope doubles by 2014, to a 1056 point increase in 90th percentile of HHI for every doubling of market size. The degree of top specialization in the largest markets rises from around 5000 to over 8000 HHI points over this time period. These patterns suggest that the right tail of specialization is highly constrained by market size, consistent with Smith (1776), and that, over time, market size have become an even more important determinant of top specialization.

In contrast, the typical oncologist faces different constraints to specialization. Figure 8 and Table 12 show that median specialization within a market is completely flat across market size in 1999. In 2014, it slopes up gently, but is still close to flat; doubling market size leads to predicted increase in median HHI of 130 points. This pattern is highly consistent

with coordination costs limiting the specialization for the typical oncologist, as in Becker and Murphy (1992). That suggests that either the returns to specialization are quite low or coordination costs are quite high for the typical oncologist.

The growing gradient in top specialization with respect to market size suggests growing economies of scale in highly specialized healthcare, with implications for patient access. Patients in large markets have increased local geographic access to sub-specialist care, while such access has stagnated or declined in smaller markets. These changes could worsen geographic differences in health outcomes if subspecialized oncologists provide higher-quality care and patient travel for care remains costly, as observed in Dingel et al. (2023).

7.2 Sources of differences in coordination costs and returns to specialization

What are the sources of coordination costs and returns to specialization, and why might they vary between specialists and generalists? Factors driving coordination costs include heterogeneous profitability across cancer types and firm-level economies of scale. Specialists may also face higher returns to specialization than generalists because of academic and organizational incentives.

One cited barrier to specialization in conversation with oncologists is that some cancers are more profitable to treat than others, disincentivizing doctors to give up patients with particularly profitable cancers. A significant portion of the revenues in oncology come from Part B dispensing fees for physician-administered drugs, where physicians receive an additional 3-6% of the average sales prices of a drug in payment for the hassle of acquiring the drug (Polite, Conti, and Ward, 2015). These fees are paid out in addition to compensation for the time spent evaluating the patient and physically administering the drug. These dispensing fees drive the notably high salaries for oncologists [13]. Empirically, the potential to

[13]To quote one oncologist, "without those [payments], we're just like any other internist', referring to the fact that oncologist salaries are among the highest for internal medicine subspecialties (Gottlieb et al., 2023) because of these payments.

earn these dispensing fees varies widely by cancer type. Appendix Figure A.8 shows that in 2014, average annual dispensing fees ranged from an estimated $45 for per testicular cancer patient to $2,175 per non-Hodgkins' lymphoma patient[14]. This finding supports the anecdotal observation that different cancers are differentially profitable and is consistent with the presence coordination costs in the form of contracting and incentive compatibility challenges.

Firm-level economies of scale may help mitigate these coordination costs for specialists but cannot entirely explain the differences. Large organizations may provide the financial and legal structures to share patients and the resultant revenue, enabling oncologists in those organizations to become specialized at baseline and in response to knowledge growth. Consistent with this explanation, specialists work in larger firms than generalists on average (Table 3); specialists work in organizations with 21 other oncologists ($\sigma = 38$) versus 17 ($\sigma = 40$) for generalists. However, these firms size distributions overlap substantially, indicating many generalists work at large firms. Firms do not entirely remove the constraints to specialization. In addition, the constraints to specialization appear to not prevent the other joint activities at the firm level (e.g., sharing of capital equipment).

An additional possibility is that specialists and generalists face different private returns to specialization. Academic institutions in particular may provide non-pecuniary incentives to become specialized. For example, oncologists notes that academic tenure and promotion systems tend to incentivize physicians to develop clinical and research expertise in a narrow domain. This story is consistent with the observation that academic physicians make up a disproportionate share of specialists (Table 3). These non-pecuniary incentives are likely weaker for most non-academic, generalist physicians, resulting in less specialization at baseline and in response to knowledge growth.

[14]Based on a dispensing fee of 6% of average sales price in the last quarter of 2014.

8 Conclusion

This paper examines how growth in general knowledge impacts expert specialization and subsequent technology adoption. I study this in the context of medical oncology, where I document large and rapid increases in anti-cancer drugs and the corpus of clinically relevant knowledge.

Descriptively, I find that they typical oncologist sees a broad range of cancers and does not become meaningfully more specialized over time. However, the right tail of extreme specialization rises over time, with the most specialized doctors developing an even narrower clinical practice.

Studying a balanced panel of oncologists, I find that initially specialized oncologists become significantly more specialized in response to knowledge growth. In contrast, generalist oncologists do not become more specialized in response to knowledge growth.

The growing dispersion in specialization has implications for patient care. Specialized oncologists prescribe newer anti-cancer drugs and have more experience with each drug they prescribe than general oncologists. In addition, these differences in drug prescribing have grown substantially over time, consistent with growing quality differences between specialist and generalist oncologists.

The growing dispersion in specialization also has impacts on the geography of healthcare. Top specialization has increased substantially in the largest markets and stagnated or declined in the smallest ones. Collectively, these results provide evidence that knowledge growth increases economies of scale and geographic inequality in specialized medical care and healthcare labor markets.

References

Acemoglu, Daron, and Joshua Linn. 2004. "Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry." *The Quarterly Journal of Economics*, 119(3): 1049–1090.

Alcacer, Juan, and Michelle Gittelman. 2006. "Patent Citations as a Measure of Knowledge Flows: The Influence of Examiner Citations." *The Review of Economics and Statistics*, 88(4): 774–779.

American Cancer Society. 2024. "Lifetime Probability of Developing (2017-2019) or Dying (2018-2020) from Cancer (Cancer Facts & Figures 2024 Supplemental Data)." Accessed on January 18, 2024.

Arrow, Kenneth J, L Kamran Bilir, and Alan Sorensen. 2020. "The impact of information technology on the diffusion of new pharmaceuticals." *American Economic Journal: Applied Economics*, 12(3): 1–39.

Becker, G. S., and K. M. Murphy. 1992. "The Division of Labor, Coordination Costs, and Knowledge." *The Quarterly Journal of Economics*, 107(4): 1137–1160.

Borusyak, Kirill, and Peter Hull. 2020. "Non-Random Exposure to Exogenous Shocks: Theory and Applications." National Bureau of Economic Research.

Budish, Eric, Benjamin N Roin, and Heidi Williams. 2015. "Do firms underinvest in long-term research? Evidence from cancer clinical trials." *American Economic Review*, 105(7): 2044–2085.

Chandra, Amitabh, David Malenka, and Jonathan Skinner. 2014. "The diffusion of new medical technology: The case of drug-eluting stents." In *Discoveries in the Economics of Aging*. 389–403. University of Chicago Press.

Dingel, Jonathan I, Joshua D Gottlieb, Maya Lozinski, and Pauline Mourot. 2023. "Market Size and Trade in Medical Services."

Duffin, Jacalyn. 2019. *Stanley's Dream: The Medical Expedition to Easter Island.* Vol. 247, McGill-Queen's Press-MQUP.

Facchini, Gabriel. 2022. "Forgetting-by-not-doing: The Case of Surgeons and Cesarean Sections." *Health Economics*, 31(3): 481–495.

for Disease Control and Prevention, Centers. 2019. "10 Leading Causes of Death, United States 2020, Both Sexes, All Ages, All Races."

Glied, Sherry, and Adriana Lleras-Muney. 2008. "Technological Innovation and Inequality in Health." *Demography*, 45(3): 741–761.

Gottlieb, Joshua D, Maria Polyakova, Kevin Rinz, Hugh Shiplett, and Victoria Udalova. 2023. "Who Values Human Capitalists' Human Capital? The Earnings and Labor Supply of U.S. Physicians." National Bureau of Economic Research.

Hockenberry, Jason M, and Lorens A Helmchen. 2014. "The Nature of Surgeon Human Capital Depreciation." *Journal of Health Economics*, 37: 70–80.

Jain, Mukesh K., Vivian G. Cheung, Paul J. Utz, Brian K. Kobilka, Tadataka Yamada, and Robert Lefkowitz. 2019. "Saving the Endangered Physician-Scientist — A Plan for Accelerating Medical Breakthroughs." *New England Journal of Medicine*, 381(5): 399–402.

Jones, Benjamin F. 2009. "The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder?" *Review of Economic Studies*, 35.

McGivney, Bill. 2008. "The NCCN Compendium for Cancer Management." *American Health and Drug Drug Benefits*, 1(5): 40–44. Interview with Bill McGivney, PhD.

Mulcahy, Nick. 2010. "First-Ever Guideline on Mesothelioma Issued by NCCN." *Medscape*. Accessed: 2024-08-06.

NCCN. 2023*a*. "About the National Comprehensive Cancer Network." *https://www.nccn.org/home/about*, Accessed on January 15, 2024.

NCCN. 2023*b*. "History of the National Comprehensive Cancer Network." *https://www.nccn.org/home/about/nccn-history*, Accessed on January 15, 2024.

Olivier, Timothée, Alyson Haslam, and Vinay Prasad. 2021. "Anticancer drugs approved by the US food and drug administration from 2009 to 2020 according to their mechanism of action." *JAMA Network Open*, 4(12): e2138793–e2138793.

Pantziarka, Pan, Rica Capistrano I, Arno De Potter, Liese Vandeborne, and Gauthier Bouche. 2021. "An open access database of licensed cancer drugs." *Frontiers in Pharmacology*, 12: 627574.

Polite, Blase, Rena M Conti, and Jeffery C Ward. 2015. "Reform of the Buy-and-Bill System for Outpatient Chemotherapy Care Is Inevitable: Perspectives from an Economist, a Realpolitik, and an Oncologist." *American Society of Clinical Oncology Educational Book*, 35(1): e75–e80.

Popescu, RA, Robert Schäfer, Raffaele Califano, Robert Eckert, Robert Coleman, J-Y Douillard, Andrés Cervantes, Paolo G Casali, Cristiana Sessa, Eric Van Cutsem, et al. 2014. "The current and future role of the medical oncologist in the professional care for cancer patients: a position paper by the European Society for Medical Oncology (ESMO)." *Annals of Oncology*, 25(1): 9–15.

Priem, Jason, Heather Piwowar, and Richard Orr. 2022. "OpenAlex: A fully-open index of scholarly works, authors, venues, institutions, and concepts." *arXiv preprint arXiv:2205.01833*.

Roach, Michael, and Wesley M Cohen. 2013. "Lens or Prism? Patent Citations as a Measure of Knowledge Flows from Public Research." *Management Science*, 59(2): 504–525.

Schlafly, Andrew, and Ronnie Sebro. 2022. "Does NIH funding differ between medical specialties? A longitudinal analysis of NIH grant data by specialty and type of grant, 2011–2020." *BMJ Open*, 12(12): e058191.

Skinner, Jonathan. 2011. "Causes and consequences of regional variations in health care." In *Handbook of Health Economics*. Vol. 2, 45–93. Elsevier.

Skinner, Jonathan, and Douglas Staiger. 2015. "Technology diffusion and productivity growth in health care." *Review of Economics and Statistics*, 97(5): 951–964.

Smith, Adam. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations.* London, United Kingdom:W. Strahan and T. Cadell.

Taylor, H. 2004. "Most physicians do not participate in clinical trials because of lack of opportunity, time, personnel support and resources." *Harris Interactive: Health Care News*, 4(9): 1–8.

Thiels, Cornelius A, Kristine T Hanson, Elizabeth B Habermann, Judy C Boughey, and Travis E Grotz. 2020. "Integrated Cancer Networks Improve Compliance with National Guidelines and Outcomes for Resectable Gastric Cancer." *Cancer*, 126(6): 1283–1294.

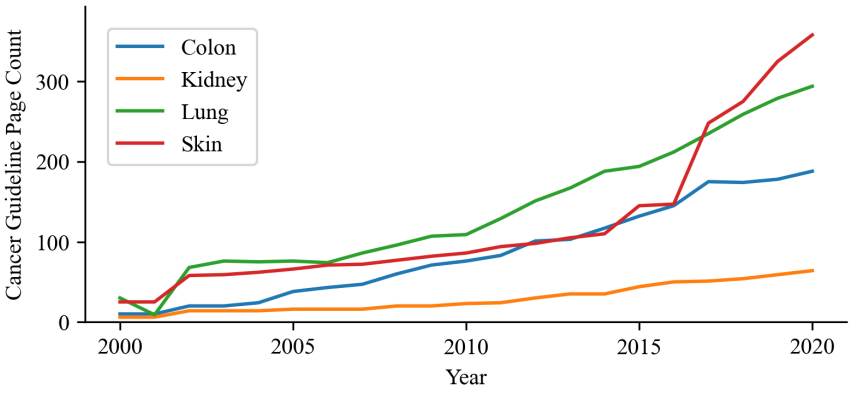
Voss, Rachel K, Yi-Ju Chiang, Keila E Torres, B Ashleigh Guadagnolo, Gary N Mann, Barry W Feig, Janice N Cormier, and Christina L Roland. 2017. "Adherence to National Comprehensive Cancer Network Guidelines is Associated with Improved Survival for Patients with Stage 2A and Stages 2B and 3 Extremity and Superficial Trunk Soft Tissue Sarcoma." *Annals of Surgical Oncology*, 24: 3271–3278.

Wallis, Claudia. 2023. "A Wave of New Cancer Treatments Challenges Community Oncologists to Keep Up." *Scientific American*. Accessed: 2024-06-07.

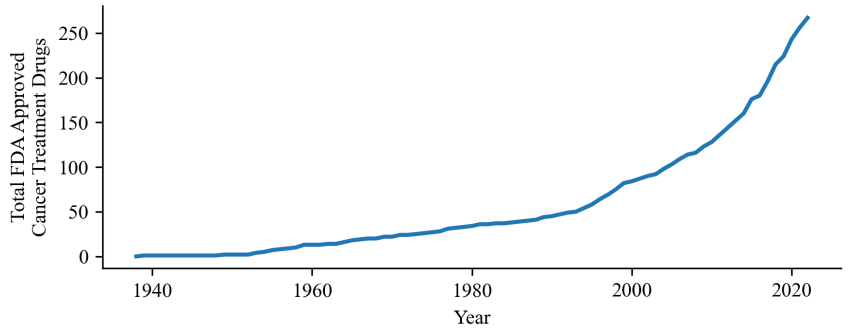
Welch, W Pete, and Andrew B Bindman. 2016. "Town and Gown Differences Among the 100 Largest Medical Groups in the United States." *Academic Medicine*, 91(7): 1007–1014.

Figure 1: Oncology has experienced large increases in knowledge and technologies

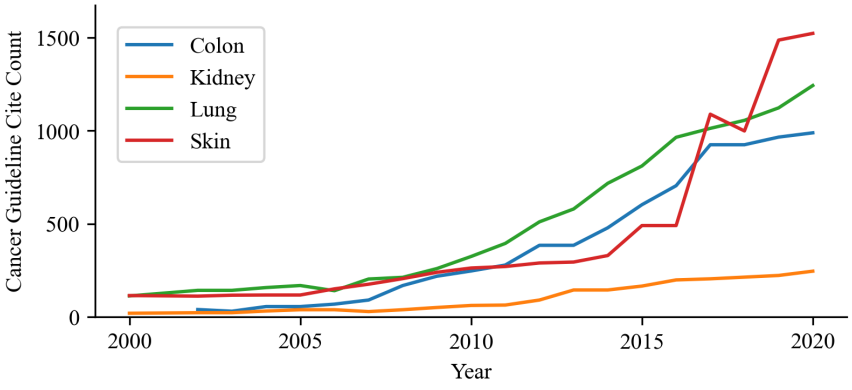
(a) Length of oncology clinical practice guidelines over time



(b) Number of FDA-approved anti-cancer drugs over time

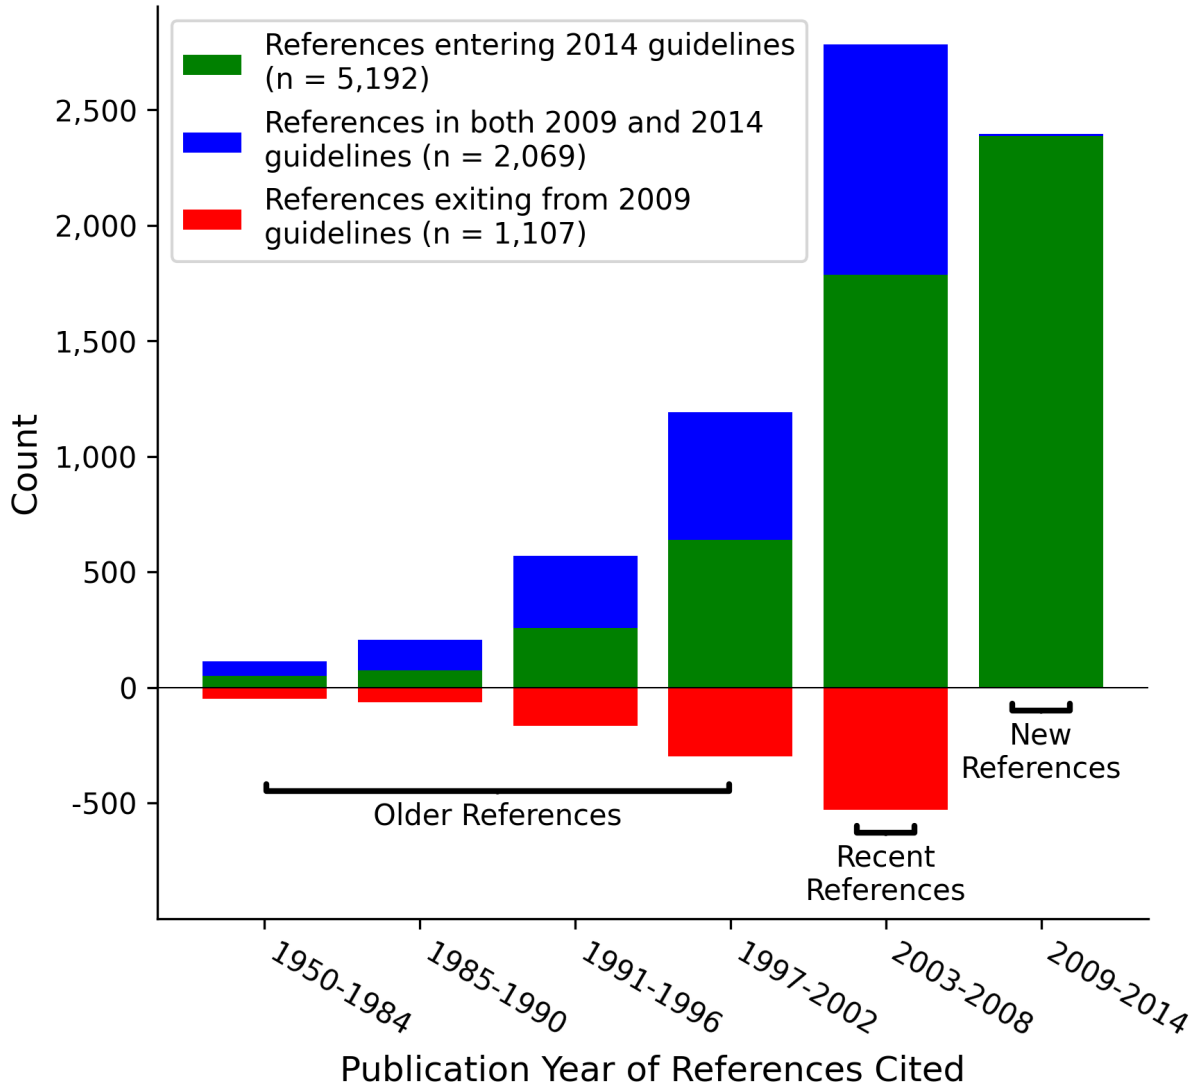


(c) Number of citations in oncology clinical practice guidelines over time



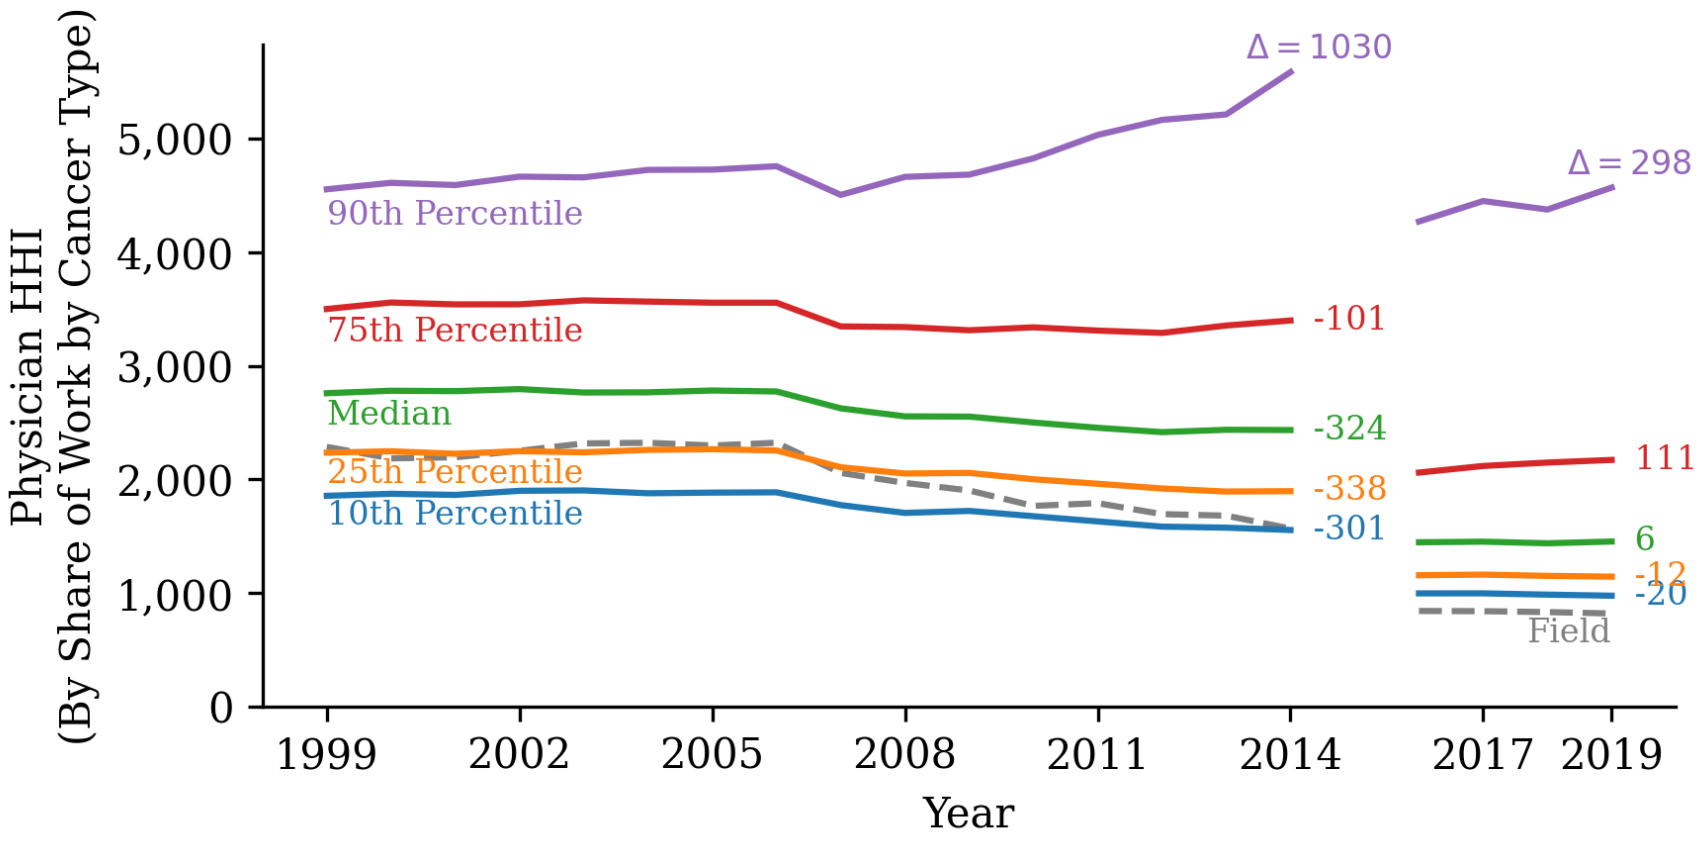
Notes: Panels (a), (b), and (c) show the growth in cancer guideline length, approved pharmaceutical treatments, and guideline citations over time. Panels (a) and (c) display the page count and citation count of the first set of treatment guidelines in each year from NCCN Guidelines. Panel (b) displays the number of unique anti-cancer drugs, where drugs are distinguished by unique Anatomical Therapeutic Chemical Level 5 Classification Codes. Note that the sample is limited to drugs with direct anti-cancer uses which are approved by the U.S. Food and Drug Administration. It does not include combinations of drugs, drugs used for cancer symptom relief or diagnosis, or unapproved drugs under investigation. The data in panel (b) is from Pantziarka et al. (2021).

Figure 2: Growth in oncology clinical practice guideline references between 2009 and 2014 comes primarily from adding references to new and recent scholarly works



Notes: Figure 2 describes changes in the population of scholarly works cited by NCCN Clinical Practice Guidelines in Oncology between the 2009 and 2014 guidelines. The x-axis shows the year bin in which the work was published. The y-axis shows the count of works. The count of works entering the 2014 guidelines are shown in green, the count of works in both 2009 and 2014 guidelines are shown in blue, and the count of works exiting the 2009 guidelines are shown in red. The sample includes all unique works in the first version of the 2009 and 2014 guidelines which can be matched with a high degree of confidence to published scholarly works in OpenAlex, an open source catalog of scholarly data (Priem, Piwowar, and Orr, 2022). This matched sample includes 83% of the citation entries found in the guidelines.

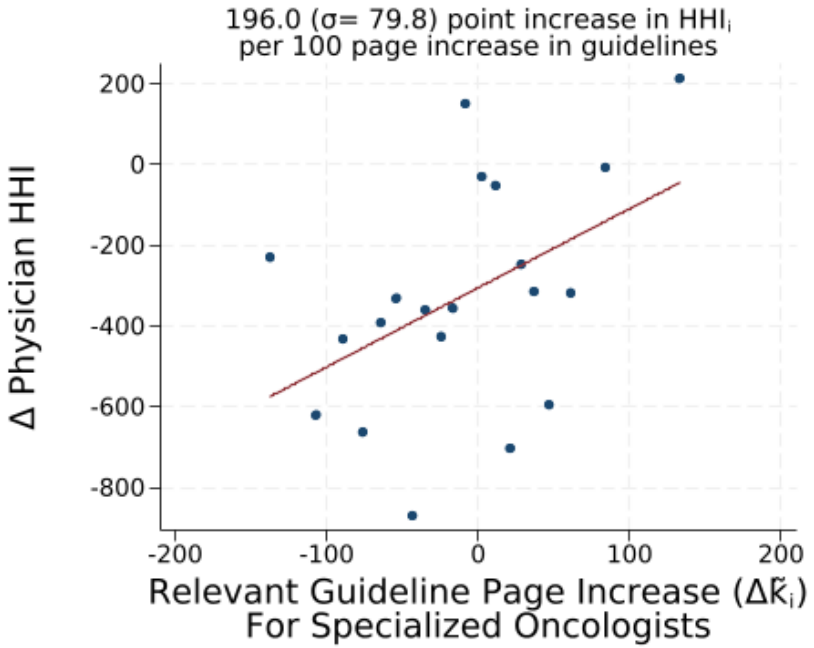
Figure 3: Specialized oncologists are growing more specialized over time



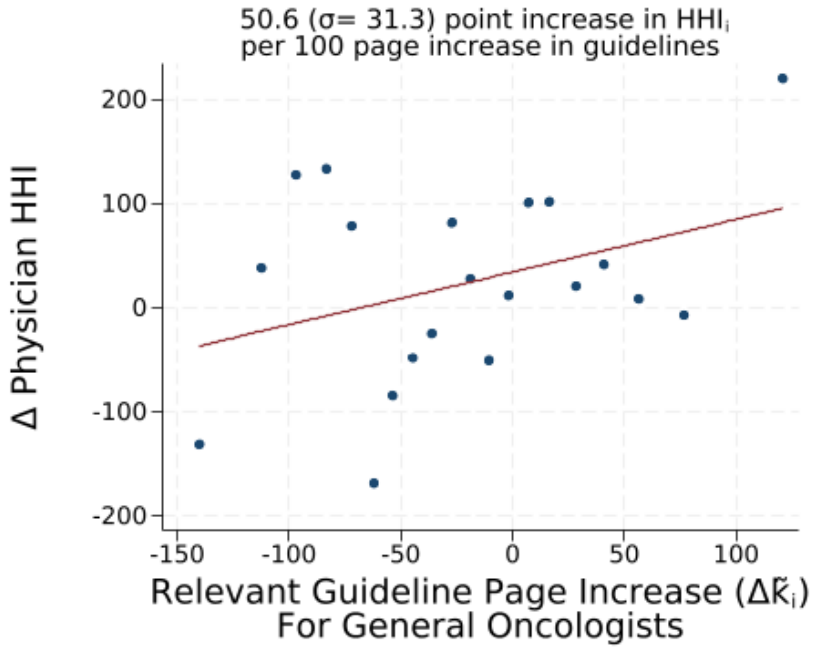
Notes: This figure shows the distribution of physician-level Herfindahl-Hirschman Index (HHI), a measure of specialization, from 1999-2019. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in the year in a 20% sample of Medicare Part B outpatient evaluation and management claims. Physician HHI is calculated within physician using shares of time spent evaluating and managing different types of cancer. Time is proxied using the work Relative Value Units for each evaluation and management HCPCS code. The 10th, 25th, 50th, 75th, and 90th percentile are displayed and labeled. The dashed grey line represents the HHI of the field overall. The trend break in 2015 is attributed to the switch in diagnosis coding schemes, from ICD-9 to ICD-10, and 2015 is dropped from the graph for visual clarity.

Figure 4: Exposure to knowledge growth increases specialization only for oncologists who start specialized (2009-2014)

(a) Specialist oncologist (top 25% of initial specialization distribution)

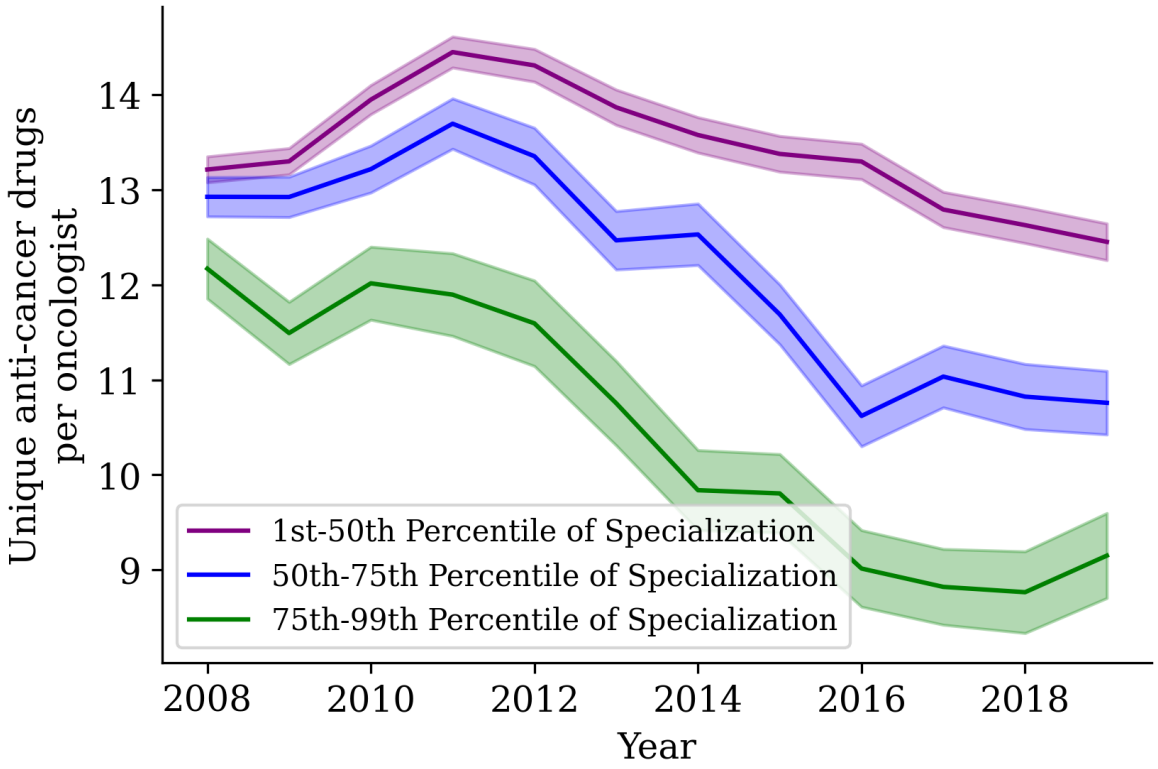


(b) Generalist oncologists (lower 75% of initial specialization distribution)



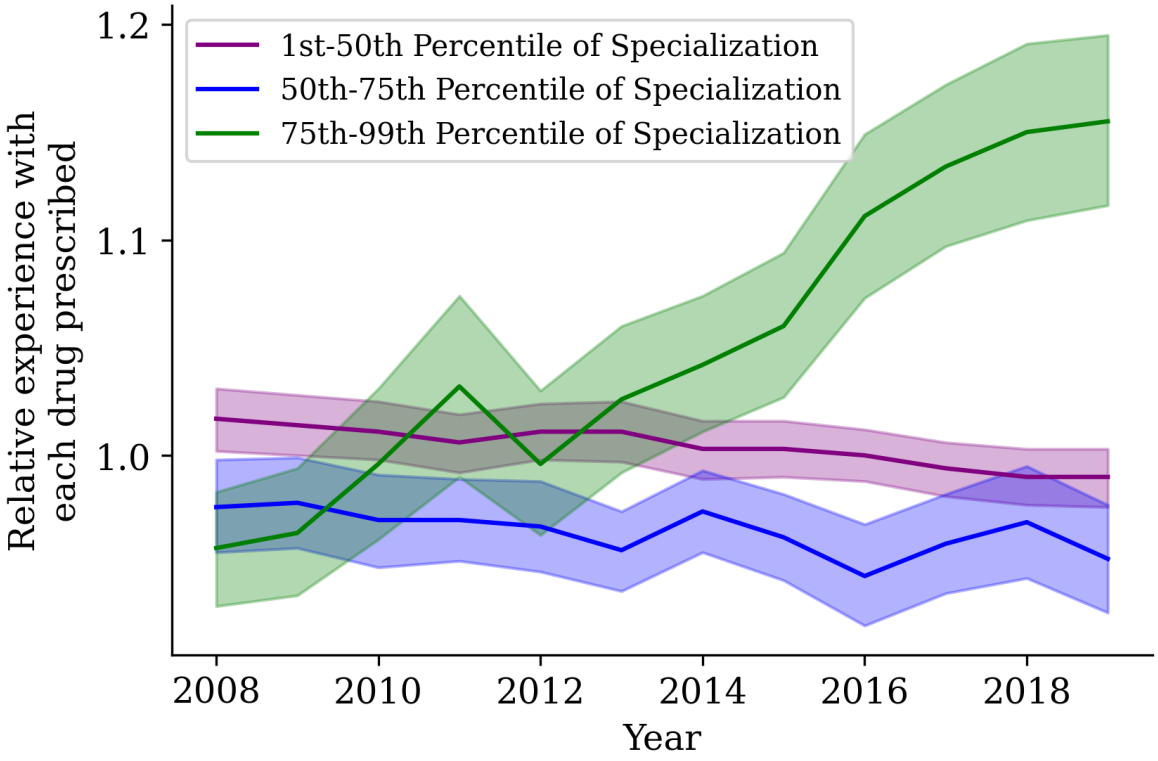
Notes: Panels (a) and (b) shows binned scatterplots and key regression estimates for the relationship between the recentered exposure to knowledge growth and change in specialization. Specialization is measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. Panel (a) shows this relationship for oncologists who start in the top 25% of the initial specialization distribution. Panel (b) shows this relationship for oncologists who start in the bottom 75% of the initial specialization distribution. In the recentered measure, the predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. Predicted exposure is estimated based on the cancer types a physician treats in 2008 and the predicted, rather than realized, growth in knowledge over the next five years. The predicted knowledge growth for each type of cancer is estimated based on a Poisson regression of the number of Medicare patients with the cancer types in 2008, an indicator for if the cancer type has well-established surrogate endpoints, and the interaction of the two. Regression standard errors are robust. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014.

Figure 5: Over time, specialists begin to prescribe fewer unique anti-cancer drugs per year than generalists



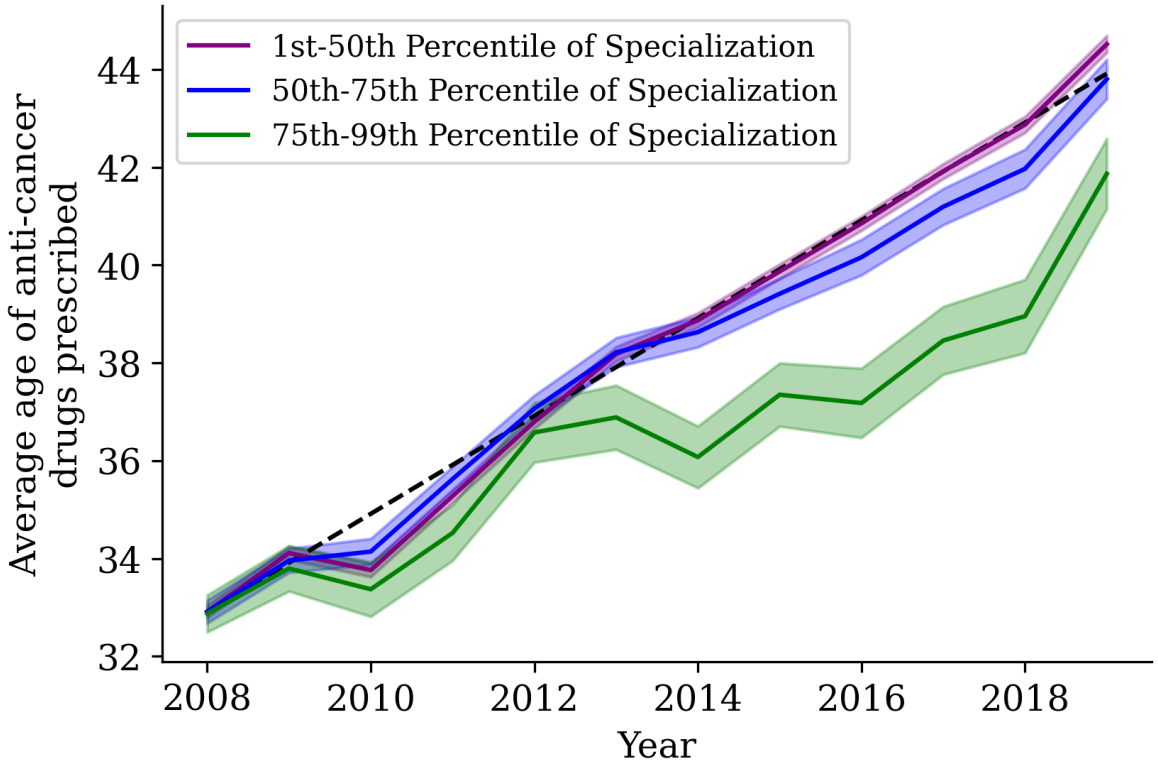
Notes: This figure shows the average number of unique anti-cancer drugs that oncologists prescribe over time by level of specialization. Unique anti-cancer drugs are defined as unique World Health Organization Anatomical Therapeutic Class codes that fall into the list of anti-cancer drugs from Pantziarka et al. (2021). Oncologists are grouped into specialization bins based on the percentile of their HHI, where a higher HHI indicates a more specialized oncologist. Within each specialization bin, the sample is restricted to oncologists who prescribed anti-cancer drugs to at least 11 patients across a 20% sample of Part B and Part D data. HHI is calculated using shares of evaluation and management work Relative Value Units over types of cancer, and the percentile is calculated among the population of oncologists with at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims.

Figure 6: Over time, specialists become relatively more experienced with the drugs they prescribe than generalists



Notes: This figure shows oncologists' average relative experience with the anti-cancer drugs that they prescribe over time by level of specialization. Experience is calculated as how many times an oncologist prescribed a drug in a year. Relative experience normalizes these values to have a mean of one in a year. Anti-cancer drugs are defined as unique World Health Organization Anatomical Therapeutic Class (ATC) codes that fall into the list of anti-cancer drugs from Pantziarka et al. (2021). Oncologists are grouped into quantile bins based on the percentile of their HHI, where a higher HHI indicates a more specialized oncologist. Within each quantile bin, the sample is restricted to oncologists who prescribed anti-cancer drugs to at least 11 patients across a 20% sample of Part B and Part D data. HHI is calculated using shares of evaluation and management work Relative Value Units over types of cancer, and the percentile is calculated among the population of oncologists with at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims.

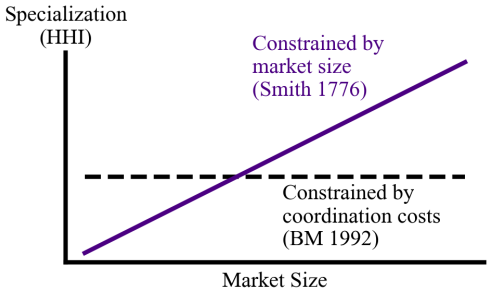
Figure 7: Over time, specialists begin to prescribe newer drugs than generalists



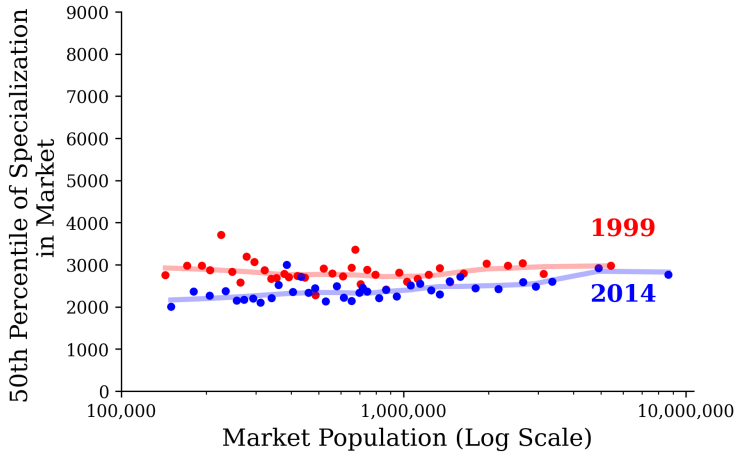
Notes: This figure shows the average age of anti-cancer drugs that oncologists prescribe over time by level of specialization. Anti-cancer drugs are defined as unique World Health Organization Anatomical Therapeutic Class (ATC) codes that fall into the list of anti-cancer drugs from Pantziarka et al. (2021). Drug age is measured as the time since FDA approval of the ATC code. Drugs are weighted by the number of patients the oncologist prescribed the drug to. Oncologists are grouped into quantile bins based on the percentile of their HHI, where a higher HHI indicates a more specialized oncologist. Within each quantile bin, the sample is restricted to oncologists who prescribed anti-cancer drugs to at least 11 patients across a 20% sample of Part B and Part D data. HHI is calculated using shares of evaluation and management work Relative Value Units over types of cancer, and the percentile is calculated among the population of oncologists with at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims.

Figure 8: Market size and specialization (1999-2014)

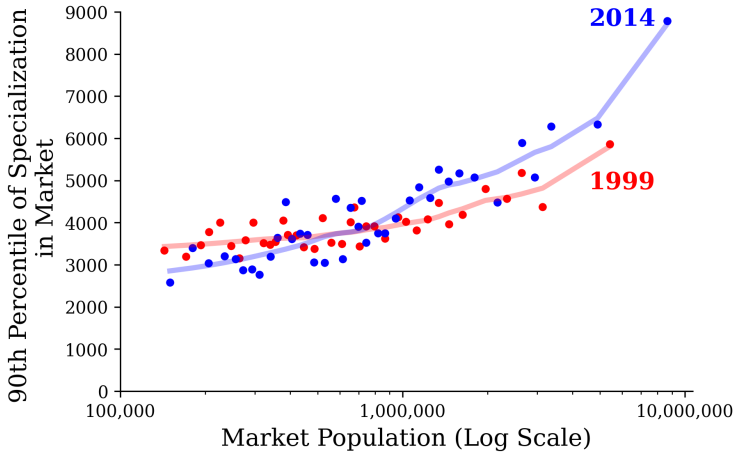
(a) Market size and coordination costs constraints lead to different patterns of specialization by market size



(b) Median specialization is stable over time and across market sizeover time



(c) 90th percentile of specialization increases more with market size



Notes: Panel (a) shows the predicted relationship between specialization and market size if specialization is constrained by market size (Smith, 1776) versus coordination costs (Becker and Murphy, 1992). Panels (b) and (c) show a binned scatterplot of the 50th and 90th percentile of oncologist specialization by market size in 1999 and 2014. The x-axis shows the market size, measured as population, in the relevant year. Markets are defined using Hospital Referral Regions (HRR). The y-axis shows oncologist specialization, as measured using a Herfindahl-Hirschman Index (HHI) over the types of cancer that the oncologist evaluated and managed in that year. In panel (b), each point represents the local average of the 50th percentile of HHI in an HRR in a year. In panel (c), each point represents the local average of the 90th percentile of HHI in an HRR in a year. The lines represent locally weighted scatterplot smoothing (LOWESS) over these points. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in the year in a 20% sample of Medicare Part B outpatient evaluation and management claims.

Table 1: Guideline length is closely correlated with the number of sources referenced in guidelines

	(1) Page Count (2009)	(2) Page Count (2014)	(3) Change in Page Count (2009-2014)
Citation count (2009)	0.354		
	(0.0252)		
Citation count (2014)		0.272	
		(0.00947)	
Change in citation count (2009-2014)			0.212
			(0.00430)
Observations	34	34	34
R-squared	0.879	0.951	0.939
Estimator	OLS	OLS	OLS
Standard errors	Robust	Robust	Robust

Notes: Table 1 shows the relationships between length of guidelines, measured in pages, and the number of sources cited in cross-section and over time. Column 1 contains a cross-sectional regression on guideline length on number of citations in 2009. Column 2 contains the same regression for 2014. Column 3 regresses changes in guideline length on changes in citation counts between 2009 and 2014. All page counts and citation counts are for the first set of guidelines in each year from NCCN Guidelines. Guidelines are grouped using the cancer grouping discussed in section 2.5 and appendix section B.2 to maintain a stable sample of guidelines over time that match with diagnosis code groupings.

Table 2: Summary statistics for balanced panel of oncologists (2008-2014)

	Mean	SD	Min	Max	N
N	3,948				
Unique Patients (2008)	61	29	25	291	3,948
Count of Cancer Types (2008)	14	4	1	25	3,948
HHI (2009)	2,920	1,517	1,028	10,000	3,948
HHI (2014)	2,855	1,660	897	10,000	3,948
Guideline Length (2009)	1,789				
Relevant Guideline Length (2009)	1,086	233	42	1,664	3,948
Change in Length of Guidelines (2009-2014)	1,745				
Change in Length of Relevant Guidelines (2009-2014)	1,032	227	24	1,606	3,948
Count Oncologists at Main Tax ID (2008)	18	40	1	219	3,948
Physician age (2008)	49	9	31	77	3,948

Notes: This table shows summary statistics for a balanced panel of oncologists from 2008 to 2014. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014. More information about each variable is calculated is detailed in section 2.

Table 3: Summary statistics for a balanced panel of specialist and generalist oncologists (2008-2014)

	Specialists (Top 25% HHI)		Generalists (Lower 75% HHI)	
	Mean	Std	Mean	Std
N	985		2,953	
Unique Patients	54	(25)	64	(30)
Count of Cancer Types	11	(4)	15	(3)
HHI (2008)	4,866	(1,866)	2,269	(485)
HHI (2009)	4,441	(2,160)	2,408	(680)
HHI (2014)	4,113	(2,453)	2,433	(974)
Academic	0.24	(0.43)	0.06	(0.23)
Share in Largest 20% of Regions	0.60	(0.49)	0.47	(0.50)
Count Oncologists at Main Tax ID	21	(38)	17	(40)
Physician Age	50	(9)	49	(9)
Relevant Guideline Length (2008)	792	(263)	997	(156)
Change in Length of Relevant Guidelines (2009-2014)	870	(289)	1,087	(170)

Notes: This table compares key summary statistics for "specialist" and "general" oncologists. Specialist oncologists are defined as those in the top 25% of the specialization distribution as measured by a Herfindahl—-Hirschman Index (HHI) over types of cancer managed in 2008. Specialist oncologists are defined as those in the bottom 75% of the specialization distribution in 2008. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014. More information about how each variable is detailed in section 2.

Table 4: Guidelines grow faster for cancer types with more patients and with high-quality surrogate endpoints

	(1) Page Count Increase	(2) Page Count Increase	(3) Page Count Increase
Log Patient Count in Medicare (2008)	0.312	0.332	0.268
	(0.0736)	(0.0632)	(0.0498)
Surrogate Endpoints		1.010	-3.149
		(0.322)	(1.413)
Surrogate x Log Pat Count (2008)			0.436
			(0.135)
Observations	35	35	35
R-squared	0.216	0.588	0.771
Estimation Method	PPML	PPML	PPML

Notes: This table shows the results of a regression of the increase in guideline pages from 2009 to 2014 on the log of the prevalence of the cancer type in Medicare in 2008, an indicator if the cancer type has high-quality surrogate endpoints, and the interaction of the two. The unit of observation is a cancer type e.g. breast cancer. The regression is estimated with Poisson maximum likelihood. Change in guideline page count is calculated using the change in page length for the first version of NCCN Guidelines in 2009 and 2014. The presence of high-quality surrogate endpoints is determined per Budish, Roin, and Williams (2015). 2008 prevalence in Medicare is calculated as the number of patients with the cancer type a 20% sample of Medicare Fee-For-Service Part B claims using ICD-9 codes. Cancer types are defined using a custom grouping of ICD-9 codes that mirrors the organization of NCCN Guidelines, as detailed in Appendix B.2.

Table 5: Physician characteristics become substantially less predictive of exposure to knowledge growth after recentering

	(1) Unadjusted exposure to knowledge growth (Δk_i)	(2) Remove physician scope variation ($\Delta \tilde{k}_i$)	(3) Remove cancer prevalence variation ($\Delta \tilde{k}_i$)	(4) Remove endpoint type variation ($\Delta \tilde{k}_i$)
Market population (millions)	-15.19	-4.145	-1.467	0.316
	(2.001)	(0.772)	(0.710)	(0.541)
Count oncologists at main org	-0.382	-0.108	0.0303	0.0278
	(0.0900)	(0.0346)	(0.0332)	(0.0258)
Physician age	-1.281	0.210	0.207	-0.103
	(0.349)	(0.150)	(0.153)	(0.120)
Patient Count	4.304	0.241	-0.276	-0.0464
	(0.118)	(0.0413)	(0.0418)	(0.0350)
Observations	3,948	3,948	3,948	3,948
R-squared	0.335	0.024	0.010	0.001
Recentered		Yes	Yes	Yes

Notes: This table shows the results of a regression of exposure to knowledge growth on observable physician characteristics. The physician characteristics include the population of the physicians' Hospital Referral Region in millions, the number of other oncologists at the Tax IDs that they work at, and physicians' patient volume. All physician characteristics are from 2008, the baseline year for the panel. In Column 1, the outcome variable is raw exposure to knowledge growth, measured as the total increase in guideline page counts for cancer types from 2009 to 2014 for cancer types the physician managed in 2008. In Columns 2-4, the outcomes are increasingly recentered measures of exposure to knowledge growth. In recentering, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. In Column 2, predicted knowledge growth is estimated as mean knowledge growth per cancer type. In Column 3, the predicted knowledge growth for each type of cancer is estimated using on a Poisson regression of the number of Medicare patients with the cancer types in 2008. Column 4 is similar, but predicted knowledge growth also adds indicator for if the cancer type has well-established surrogate endpoints and the interaction between the number of Medicare patients and surrogate endpoints. The full recentering procedure removes variation in exposure to knowledge growth that arises become some physicians have a broader scope and because some cancer types predictably have more innovation. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014. All regression standard errors are robust.

Table 6: Effect knowledge growth on specialization for ex-ante specialized oncologists

	(1) Knowledge Growth Recentered ($\Delta\tilde{k}_i$)	(2) Restrict Sample	(3) Add Age FE	(4) Add Market FE	(5) Add Pat Count FE	(6) Add Org Size FE
Increase in Guideline Length (100 Pages)	196.0	213.2	220.0	253.1	255.2	246.3
	(79.81)	(81.06)	(81.56)	(89.07)	(89.98)	(91.20)
Observations	988	924	924	924	924	924
R-squared	0.007	0.008	0.009	0.172	0.177	0.178
Age FE			Yes	Yes	Yes	Yes
Market FE				Yes	Yes	Yes
Org Size FE					Yes	Yes
Pat Count FE						Yes

Notes: This table shows regression estimates for the relationship between exposure to knowledge growth and change in specialization for oncologists who start in the top 25% of the initial specialization distribution. Specialization is measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. The measure of exposure to knowledge growth is recentered to control for variation in physician scope and field expected knowledge growth. Column 1 shows the coefficient estimate for the univariate regression. The regression standard errors are robust. Column 2 restricts the sample to physicians where we observe age, hospital referral region (HRR) and the number of other oncologists at their Tax IDs. Columns 3, 4, and 5 add fixed effects for the physicians age group, hospital referral region (HRR) and the number of other oncologists at their Tax IDs, respectively. In the recentered measure of exposure to knowledge growth, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. Raw exposure to knowledge growth is calculated as the total increase in the page count of NCCN Guidelines over 2009 to 2014 for the types of cancers a physician treats in 2008. Predicted exposure is estimated based on the cancer types a physician treats in 2008 and the predicted, rather than realized, growth in knowledge over the next five years. The predicted knowledge growth for each type of cancer is estimated based on a Poisson regression of the number of Medicare patients with the cancer types in 2008, an indicator for if the cancer type has well-established surrogate endpoints, and the interaction of the two. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014.

Table 7: Effect knowledge growth on specialization for ex-ante general oncologists

	(1) Knowledge Growth Recentered ($\Delta\tilde{k}_i$)	(2) Restrict Sample	(3) Add Age FE	(4) Add Market FE	(5) Add Pat Count FE	(6) Add Org Size FE
Increase in Guideline Length (100 Pages)	50.57	54.23	54.18	66.96	66.50	66.26
	(31.32)	(31.45)	(31.44)	(33.14)	(33.13)	(33.03)
Observations	2,962	2,937	2,937	2,937	2,937	2,937
R-squared	0.001	0.001	0.001	0.108	0.108	0.110
Age FE			Yes	Yes	Yes	Yes
Market FE				Yes	Yes	Yes
Org Size FE					Yes	Yes
Pat Count FE						Yes

Notes: This table shows regression estimates for the relationship between exposure to knowledge growth and change in specialization for oncologists who start in the bottom 75% of the initial specialization distribution. Specialization is measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. The measure of exposure to knowledge growth is recentered to control for variation in physician scope and field expected knowledge growth. Column 1 shows the coefficient estimate for the univariate regression. The regression standard errors are robust. Column 2 restricts the sample to physicians where we observe age, hospital referral region (HRR) and the number of other oncologists at their Tax IDs. Columns 3, 4, and 5 add fixed effects for the physicians age group, hospital referral region (HRR) and the number of other oncologists at their Tax IDs, respectively. In the recentered measure of exposure to knowledge growth, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. Raw exposure to knowledge growth is calculated as the total increase in the page count of NCCN Guidelines over 2009 to 2014 for the types of cancers a physician treats in 2008. Predicted exposure is estimated based on the cancer types a physician treats in 2008 and the predicted, rather than realized, growth in knowledge over the next five years. The predicted knowledge growth for each type of cancer is estimated based on a Poisson regression of the number of Medicare patients with the cancer types in 2008, an indicator for if the cancer type has well-established surrogate endpoints, and the interaction of the two. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014.

Table 8: Effect knowledge growth on count of unique cancer types treated for ex-ante specialized oncologists

	(1) Knowledge Growth Recentered ($\Delta\tilde{k}_i$)	(2) Restrict Sample	(3) Add Age FE	(4) Add Market FE	(5) Add Pat Count FE	(6) Add Org Size FE
Increase in Guideline Length (100 Pages)	-0.111	-0.135	-0.172	-0.197	-0.222	-0.248
	(0.169)	(0.174)	(0.174)	(0.188)	(0.188)	(0.190)
Observations	988	924	924	924	924	924
R-squared	0.000	0.001	0.009	0.221	0.242	0.246
Age FE			Yes	Yes	Yes	Yes
Market FE				Yes	Yes	Yes
Pat Count FE					Yes	Yes
Org Size FE						Yes

Notes: This table shows regression estimates for the relationship between exposure to knowledge growth and the number of cancers managed for specialized oncologists. Specialized oncologists are defined as oncologists in the top 25% of the specialization distribution in 2009, measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. The measure of exposure to knowledge growth is recentered to control for variation in physician scope and field expected knowledge growth. Column 1 shows the coefficient estimate for the univariate regression. The regression standard errors are robust. Column 2 restricts the sample to physicians where we observe age, hospital referral region (HRR) and the number of other oncologists at the Tax IDs. Columns 3, 4, and 5 add fixed effects for the physicians age group, HRR and the number of other oncologists at their Tax IDs, respectively. In the recentered measure of exposure to knowledge growth, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. Raw exposure to knowledge growth is calculated as the total increase in the page count of NCCN Guidelines over 2009 to 2014 for the types of cancers a physician treats in 2008. Predicted exposure is estimated based on the cancer types a physician treats in 2008 and the predicted, rather than realized, growth in knowledge over the next five years. The predicted knowledge growth for each type of cancer is estimated based on a Poisson regression of the number of Medicare patients with the cancer types in 2008, an indicator for if the cancer type has well-established surrogate endpoints, and the interaction of the two. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014.

Table 9: Effect knowledge growth on count of unique cancer types treated for ex-ante general oncologists

	(1) Knowledge Growth Recentered ($\Delta\tilde{k}_i$)	(2) Restrict Sample	(3) Add Age FE	(4) Add Market FE	(5) Add Pat Count FE	(6) Add Org Size FE
Increase in Guideline Length (100 Pages)	-0.165	-0.182	-0.180	-0.236	-0.244	-0.236
	(0.0909)	(0.0909)	(0.0908)	(0.0935)	(0.0937)	(0.0935)
Observations	2,962	2,937	2,937	2,937	2,937	2,937
R-squared	0.001	0.001	0.022	0.160	0.178	0.180
Age FE			Yes	Yes	Yes	Yes
Market FE				Yes	Yes	Yes
Pat Count FE					Yes	Yes
Org Size FE						Yes

Notes: This table shows regression estimates for the relationship between exposure to knowledge growth and the number of cancers managed for general oncologists. General oncologists are defined as oncologists in the bottom 75% of the specialization distribution in 2009, measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. The measure of exposure to knowledge growth is recentered to control for variation in physician scope and field expected knowledge growth. Column 1 shows the coefficient estimate for the univariate regression. The regression standard errors are robust. Column 2 restricts the sample to physicians where we observe age, hospital referral region (HRR) and the number of other oncologists at their Tax IDs. Columns 3, 4, and 5 add fixed effects for the physicians age group, hospital referral region (HRR) and the number of other oncologists at their Tax IDs, respectively. In the recentered measure of exposure to knowledge growth, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. Raw exposure to knowledge growth is calculated as the total increase in the page count of NCCN Guidelines over 2009 to 2014 for the types of cancers a physician treats in 2008. Predicted exposure is estimated based on the cancer types a physician treats in 2008 and the predicted, rather than realized, growth in knowledge over the next five years. The predicted knowledge growth for each type of cancer is estimated based on a Poisson regression of the number of Medicare patients with the cancer types in 2008, an indicator for if the cancer type has well-established surrogate endpoints, and the interaction of the two. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014.

Table 10: The effect knowledge growth on specialization for ex-ante specialized oncologists with alternative measures of knowledge growth

	(1)	(2)	(3)	(4)
Increase in Guideline Length (100 Pages)	196.0	186.0	310.7	-20.22
	(79.81)	(86.76)	(109.5)	(162.5)
Observations	988	988	988	326
R-squared	0.007	0.004	0.008	0.000
Min Pat Count per Field	1	2	3	1
HHI Sample	All pat	All pat	All pat	All pat
Claims	Outpatient EM	Outpatient EM	Outpatient EM	New EM
Min Pat Count per Doc		25	25	11

Notes: This table shows regression estimates for the relationship between exposure to knowledge growth and change in specialization for oncologists who start in the top 25% of the initial specialization distribution. Specialization is measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. The measure of exposure to knowledge growth is recentered to control for variation in physician scope and field expected knowledge growth. Column 1 shows the baseline univariate regression. Column 2 and 3 raise the minimum patient count to be considered to be exposed to a cancer type in the baseline year. Column 4 restricts the sample to new evaluation and management claims (versus all evaluation and management claims) and lowers the threshold for a physician to be included in the sample to 11 patients per year. In the recentered measure of exposure to knowledge growth, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. Raw exposure to knowledge growth is calculated as the total increase in the page count of NCCN Guidelines over 2009 to 2014 for the types of cancers a physician treats in 2008. Predicted exposure is estimated based on the cancer types a physician treats in 2008 and the predicted, rather than realized, growth in knowledge over the next five years. The predicted knowledge growth for each type of cancer is estimated based on a Poisson regression of the number of Medicare patients with the cancer types in 2008, an indicator for if the cancer type has well-established surrogate endpoints, and the interaction of the two. Unless otherwise specified, all physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014.

Table 11: The effect knowledge growth on specialization for ex-ante general oncologists with alternative measures of knowledge growth

	(1)	(2)	(3)	(4)
Increase in Guideline Length (100 Pages)	50.57	7.200	8.637	68.64
	(31.32)	(32.01)	(34.67)	(67.85)
Observations	2,962	2,962	2,962	978
R-squared	0.001	0.000	0.000	0.001
Min Pat Count per Field	1	2	3	1
HHI Sample	All pat	All pat	All pat	All pat
Claims	Outpatient EM	Outpatient EM	Outpatient EM	New EM
Min Pat Count per Doc		25	25	11

Notes: This table shows regression estimates for the relationship between exposure to knowledge growth and change in specialization for oncologists who start in the lower 75% of the initial specialization distribution. Specialization is measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. The measure of exposure to knowledge growth is recentered to control for variation in physician scope and field expected knowledge growth. Column 1 shows the baseline univariate regression. Column 2 and 3 raise the minimum patient count to be considered to be exposed to a cancer type in the baseline year. Column 4 restricts the sample to new evaluation and management claims (versus all evaluation and management claims) and lowers the threshold for a physician to be included in the sample to 11 patients per year. In the recentered measure of exposure to knowledge growth, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. Raw exposure to knowledge growth is calculated as the total increase in the page count of NCCN Guidelines over 2009 to 2014 for the types of cancers a physician treats in 2008. Predicted exposure is estimated based on the cancer types a physician treats in 2008 and the predicted, rather than realized, growth in knowledge over the next five years. The predicted knowledge growth for each type of cancer is estimated based on a Poisson regression of the number of Medicare patients with the cancer types in 2008, an indicator for if the cancer type has well-established surrogate endpoints, and the interaction of the two. Unless otherwise specified, all physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014.

Table 12: Median and 90th percentile specialization by market size over time

	(1) Median HHI	(2) 90th Pct HHI
Current Population (Log)	-23.64	495.1
	(45.01)	(79.34)
Current Population (Log) x 2014	153.3	561.2
	(56.65)	(122.7)
Observations	601	601
R-squared	0.136	0.271

Notes: Table 12 shows a regression of market-level oncologist specialization on market size, year, and an interaction of market size and year. In Column 1, the outcome is the 50th percentile of oncologist specialization in a market. In Column 2, the outcome is the 90th percentile of oncologist specialization in a market. Markets are defined using Hospital Referral Regions (HRR). Market size is measured as population in the relevant year. Oncologist specialization is measured using a Herfindahl-Hirschman Index (HHI) over the types of cancer that the oncologist evaluated and managed in that year. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in the year in a 20% sample of Medicare Part B outpatient evaluation and management claims data.

9 Appendix

For Online Publication

Knowledge Growth and Specialization: Evidence from Oncologists

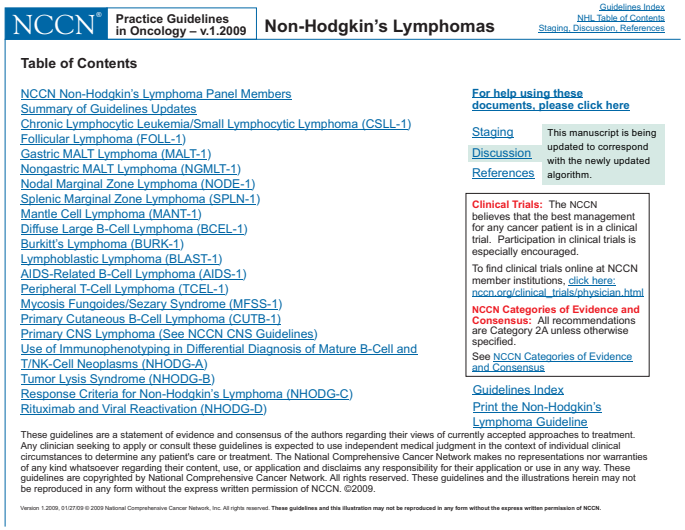
Maya Lozinski

November 2024

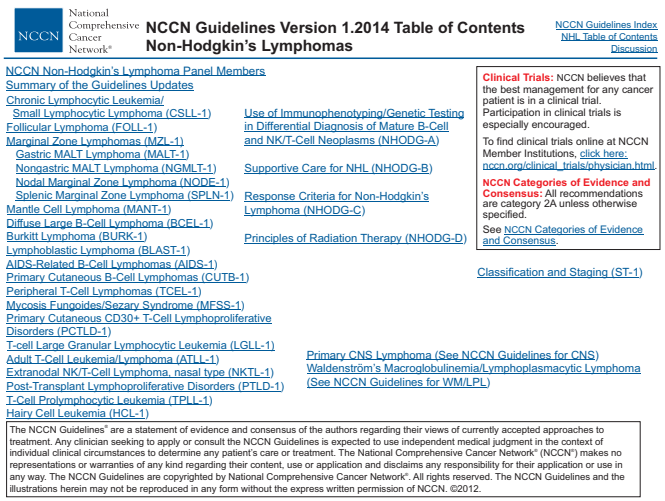
A Additional exhibits

Figure A.1: Example changes in NCCN Guidelines: Table of Contents for Non-Hodgkins Lymphoma

(a) 2009



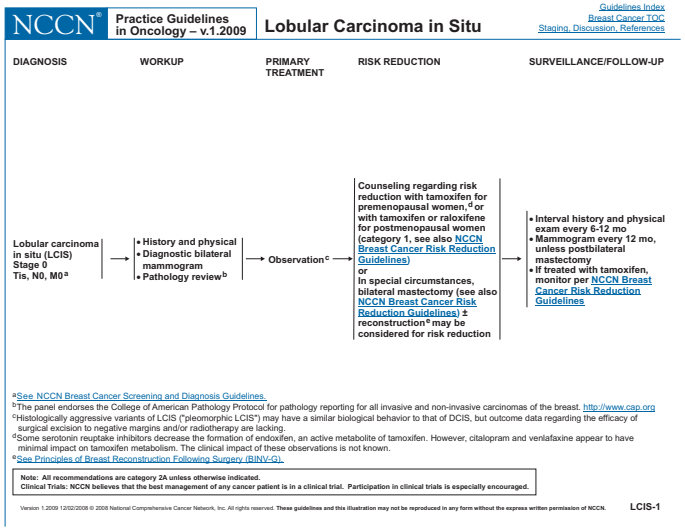
(b) 2014



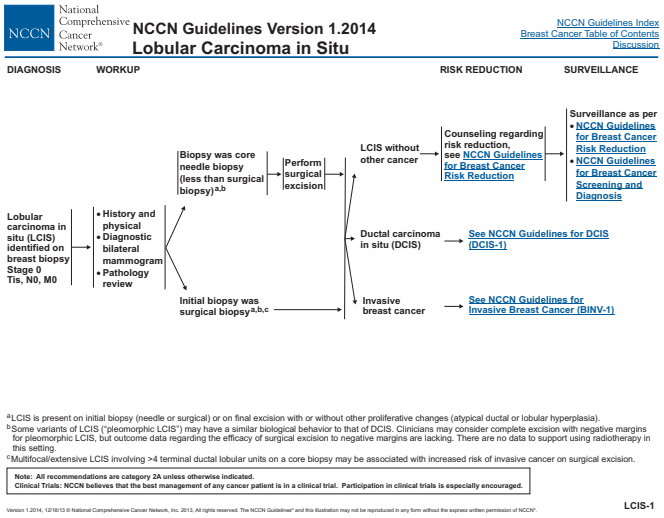
Notes: Figure A.1 shows example changes from NCCN Guidelines. Figures A.1a and A.1b show the table of contents for non-Hodgkins lymphoma (NHL) in 2009 and 2014, respectively. Figure A.1a reproduced with permission from the NCCN Clinical Practice Guidelines in Oncology (NCCN Guidelines®) for Non-Hodgkin's Lymphomas V.1.2009 © 2009 National Comprehensive Cancer Network, Inc. All rights reserved. The NCCN Guidelines® and illustrations herein may not be reproduced in any form for any purpose without the express written permission of NCCN. To view the most recent and complete version of the NCCN Guidelines, go online to NCCN.org. The NCCN Guidelines are a work in progress that may be refined as often as new significant data becomes available. Figure A.1b reproduced with permission from the NCCN Clinical Practice Guidelines in Oncology (NCCN Guidelines®) for Non-Hodgkin's Lymphomas V.1.2014. © 2013 National Comprehensive Cancer Network, Inc. All rights reserved. The NCCN Guidelines® and illustrations herein may not be reproduced in any form for any purpose without the express written permission of NCCN. To view the most recent and complete version of the NCCN Guidelines, go online to NCCN.org. The NCCN Guidelines are a work in progress that may be refined as often as new significant data becomes available. NCCN makes no warranties of any kind whatsoever regarding their content, use or application and disclaims any responsibility for their application or use in any way.

Figure A.2: Example changes in NCCN Guidelines: Flowchart for Lobular Carcinoma In-Situ

(a) 2009

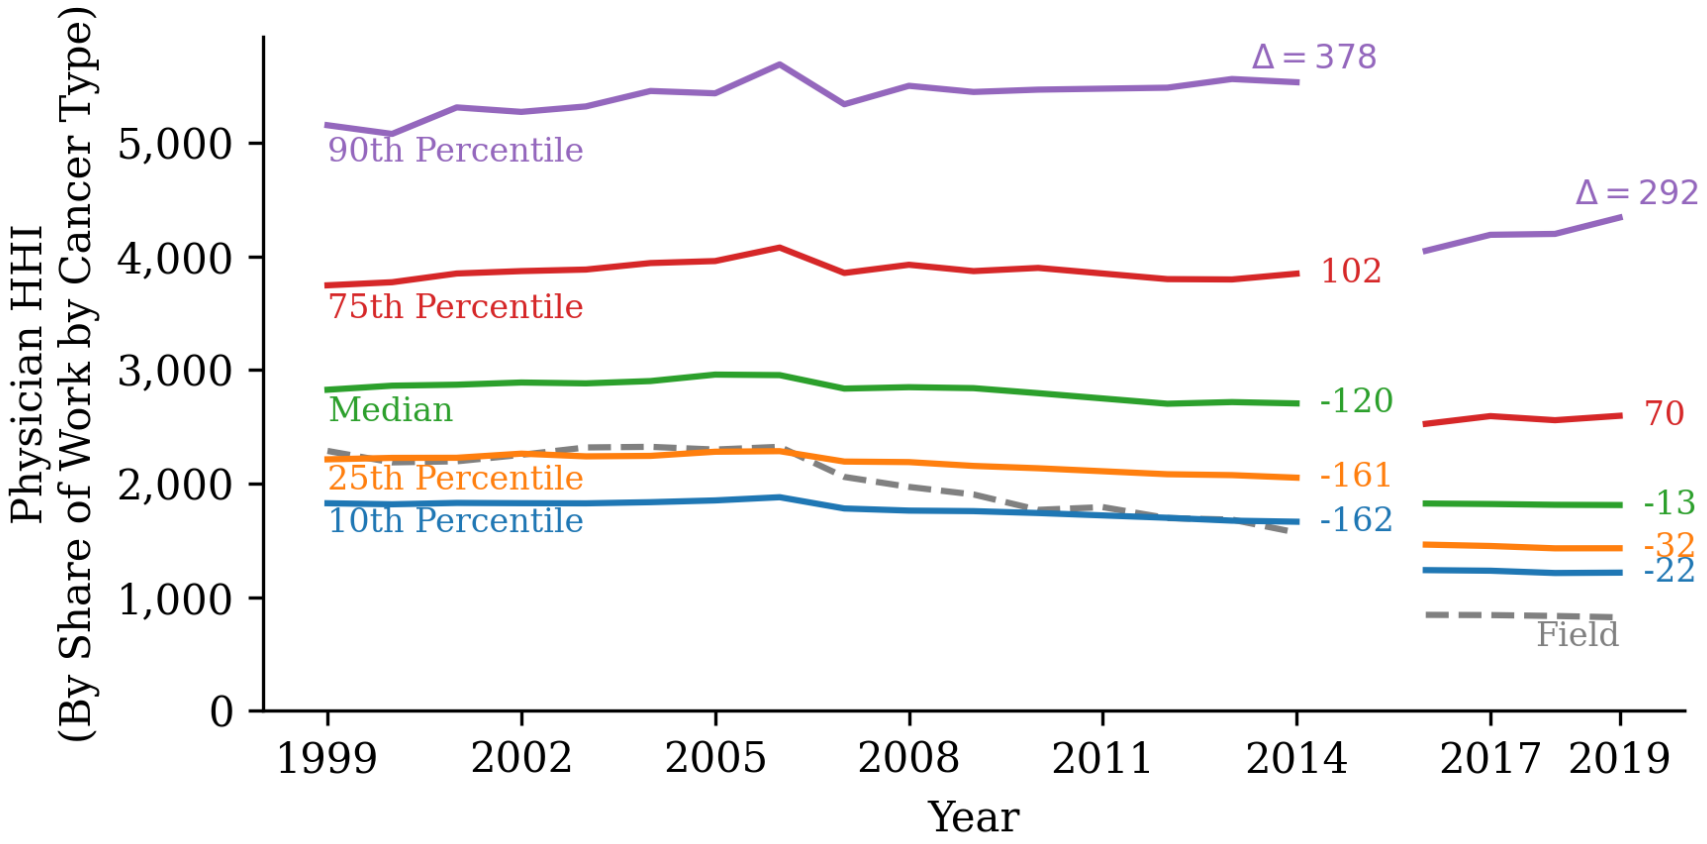


(b) 2014



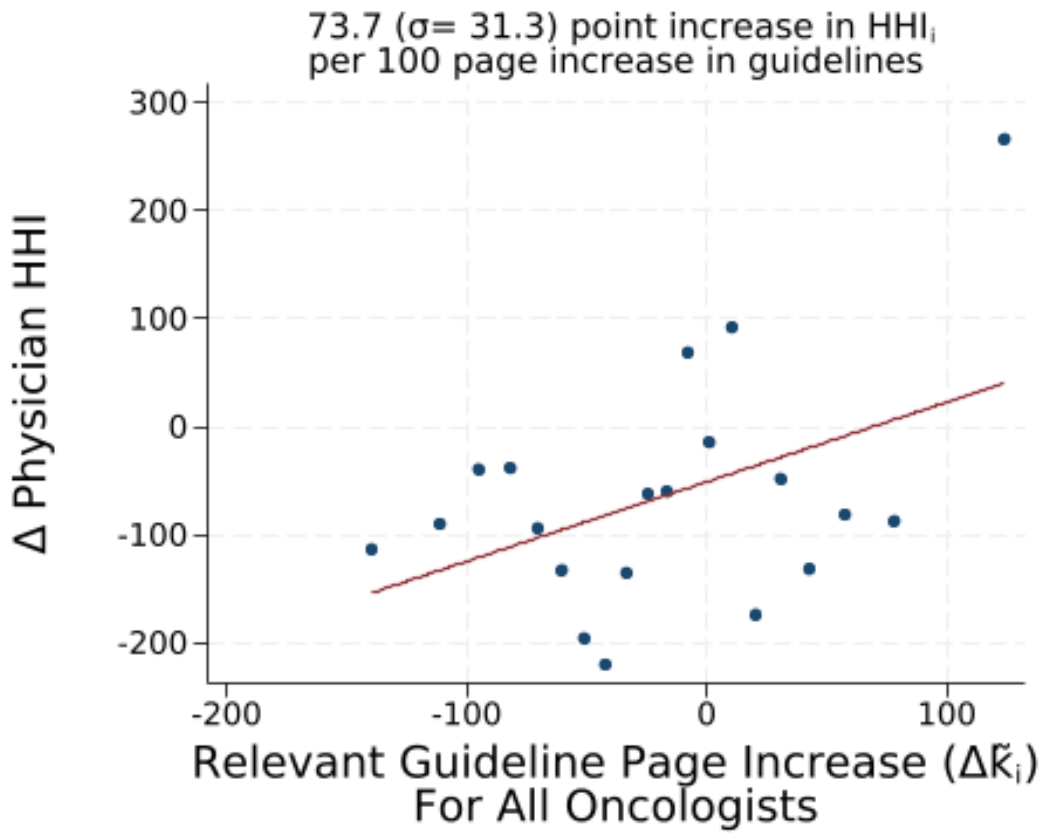
Notes: Figure A.2 shows example changes from NCCN Guidelines. Figures A.2a and A.2b show a flowchart with the diagnostic workup for lobular carcinoma in-situ, a subtype of breast cancer, in 2009 and 2014, respectively. Figure A.2a reproduced with permission from the NCCN Clinical Practice Guidelines in Oncology (NCCN Guidelines®) for Breast Cancer V.1.2009. © 2008 National Comprehensive Cancer Network, Inc. All rights reserved. The NCCN Guidelines® and illustrations herein may not be reproduced in any form for any purpose without the express written permission of NCCN. To view the most recent and complete version of the NCCN Guidelines, go online to NCCN.org. The NCCN Guidelines are a work in progress that may be refined as often as new significant data becomes available. Figure A.2b Reproduced with permission from the NCCN Clinical Practice Guidelines in Oncology (NCCN Guidelines®) for Breast Cancer V.1.2014. © 2013 National Comprehensive Cancer Network, Inc. All rights reserved. The NCCN Guidelines® and illustrations herein may not be reproduced in any form for any purpose without the express written permission of NCCN. To view the most recent and complete version of the NCCN Guidelines, go online to NCCN.org. The NCCN Guidelines are a work in progress that may be refined as often as new significant data becomes available. NCCN makes no warranties of any kind whatsoever regarding their content, use or application and disclaims any responsibility for their application or use in any way.

Figure A.3: Specialized oncologists are growing more specialized over time (25 patients used to calculate HHI)



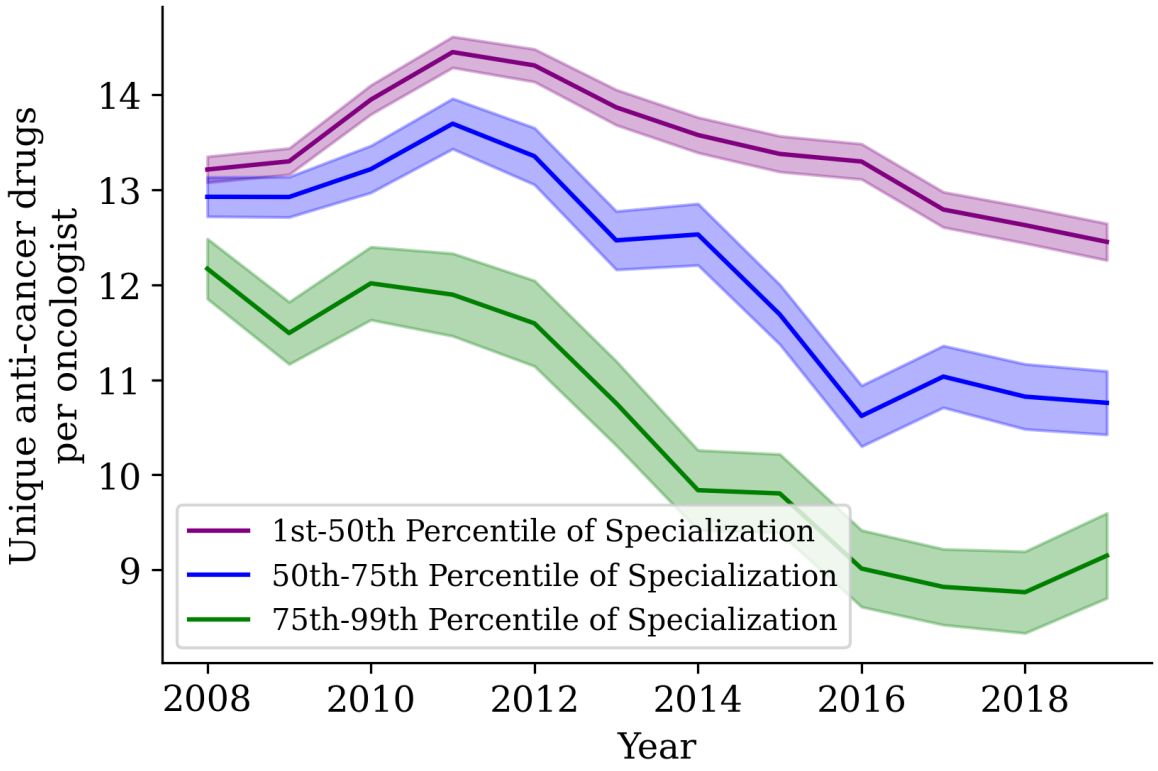
Notes: This figure shows the distribution of physician-level Herfindahl-Hirschman Index (HHI), a measure of specialization, from 1999-2019, using only 25 random patients per physician per year to calculate HHI. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in the year in a 20% sample of Medicare Part B outpatient evaluation and management claims. Physician HHI is calculated within physician using shares of time spent evaluating and managing different types of cancer. Time is proxied using the work Relative Value Units for each evaluation and management HCPCS code. The 10th, 25th, 50th, 75th, and 90th percentile are displayed and labeled. The dashed grey line represents the HHI of the field overall. The trend break in 2015 is attributed to the switch in diagnosis coding schemes, from ICD-9 to ICD-10, and 2015 is dropped from the graph for visual clarity.

Figure A.4: Effect of knowledge growth on specialization for all oncologists



Notes: Figure A.4 shows a binned scatterplot and key regression estimates for the relationship between the recentered exposure to knowledge growth and change in specialization for all oncologists. In the recentered measure, the predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. Predicted exposure is estimated based on the cancer types a physician treats in 2008 and the predicted, rather than realized, growth in knowledge over the next five years. The predicted knowledge growth for each type of cancer is estimated based on a Poisson regression of the number of Medicare patients with the cancer types in 2008, an indicator for if the cancer type has well-established surrogate endpoints, and the interaction of the two. Regression standard errors are robust. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014.

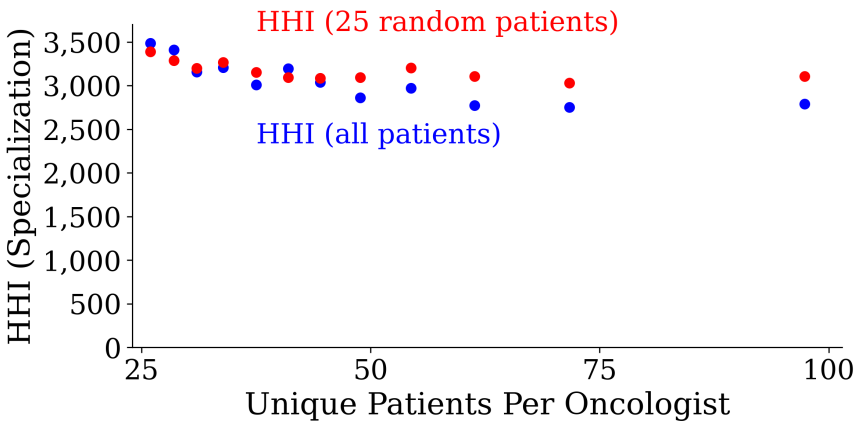
Figure A.5: Over time, specialists prescribe fewer unique anti-cancer drugs than generalists in a year, even after removing variation in patient volume



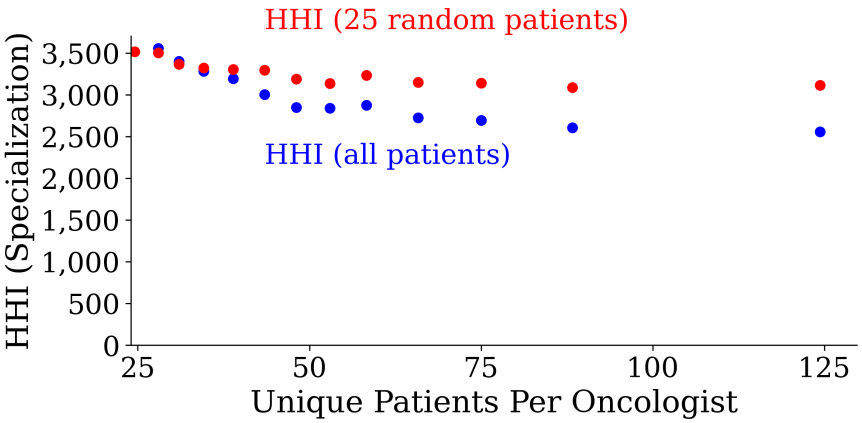
Notes: This figure shows the average number of unique anti-cancer drugs that oncologists prescribe over time by level of specialization. This figure differs from Figure 5 in that the average number of unique anti-cancer drugs is calculated over 11 random patients rather than all patients per oncologist per year. This change in sample removes systematic differences in patient volume that may exist by level of oncologist specialization. Unique anti-cancer drugs are defined as unique World Health Organization Anatomical Therapeutic Class codes that fall into the list of anti-cancer drugs from Pantziarka et al. (2021). Oncologists are grouped into specialization bins based on the percentile of their HHI, where a higher HHI indicates a more specialized oncologist. Within each specialization bin, the sample is restricted to oncologists who prescribed anti-cancer drugs to at least 11 patients across a 20% sample of Part B and Part D data. HHI is calculated using shares of evaluation and management work Relative Value Units over types of cancer, and the percentile is calculated among the population of oncologists with at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims.

Figure A.6: Physician HHI declines with patient volume

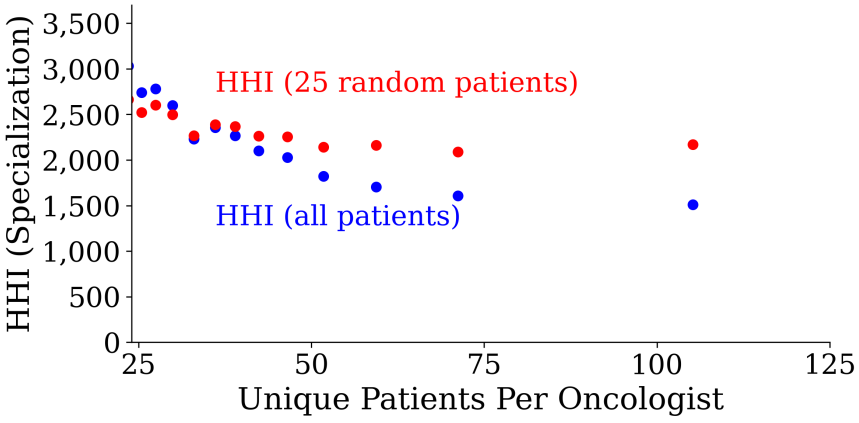
(a) 1999



(b) 2009

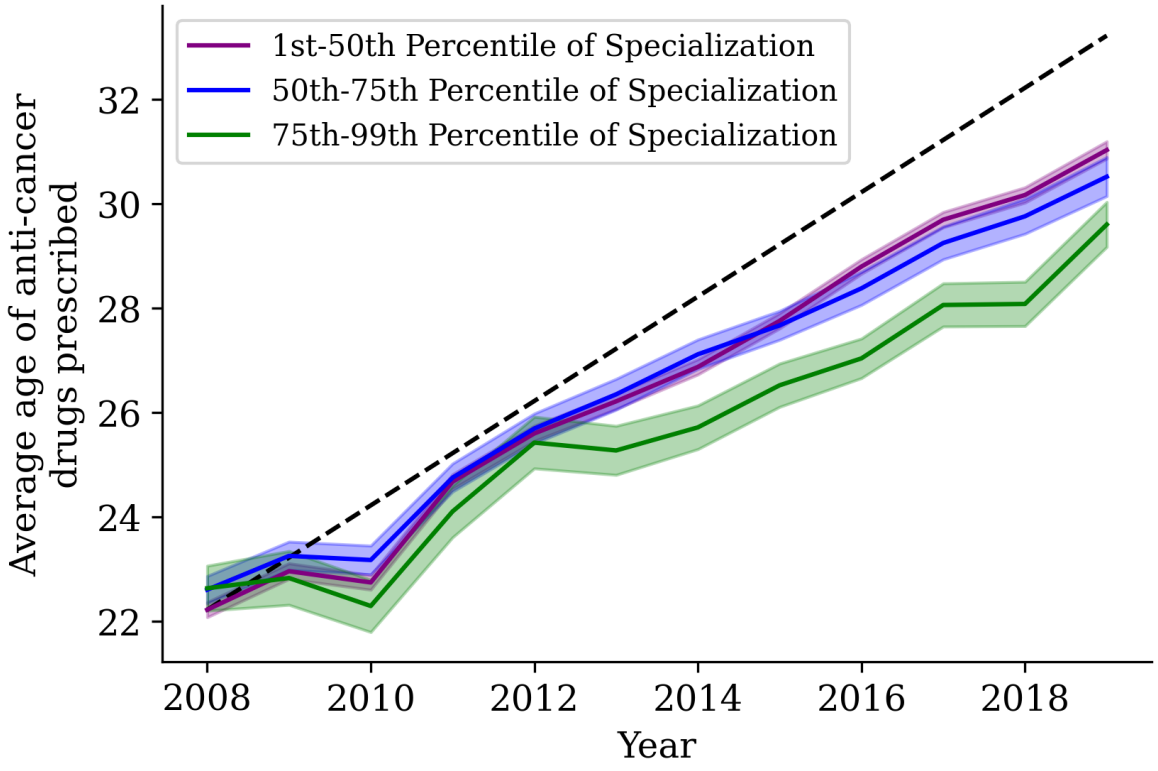


(c) 2019



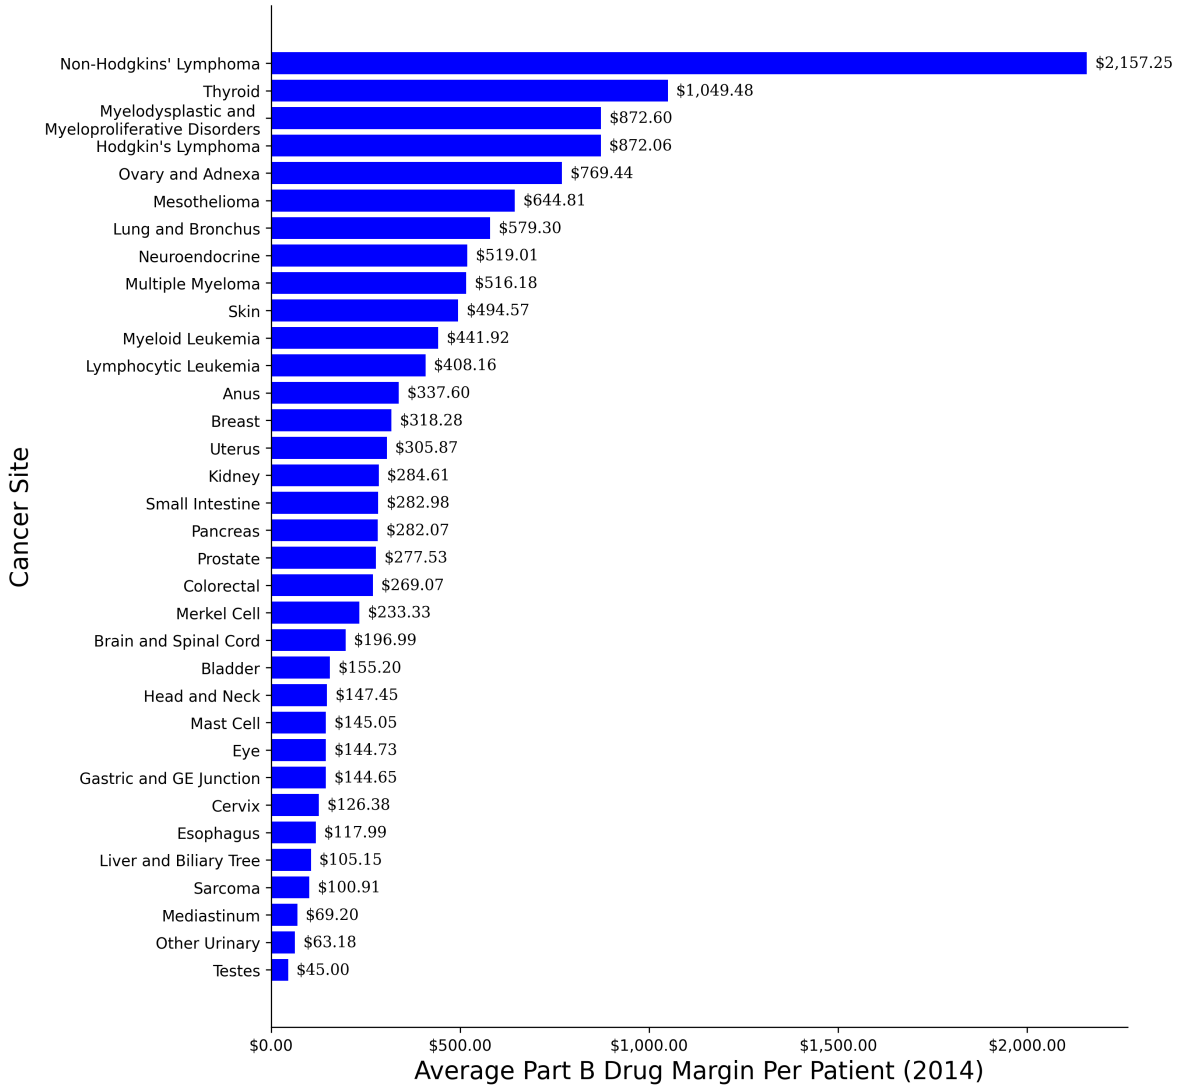
Notes: Figure displays binned scatter plots of physician-level Herfindahl-Hirschman Index (HHI), calculated using both all patients and 25 random patients per physician, by patient volume. Figures , , show the relationship for 1999, 2009 and 2019 respectively. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in the year in a 20% sample of Medicare Part B outpatient evaluation and management claims. Physician HHI is calculated within physician using shares of time spent evaluating and managing different types of cancer. Time is proxied using the work Relative Value Units for each evaluation and management HCPCS code.

Figure A.7: Over time, specialists begin to prescribe newer drugs than generalists (excluding steroids)



Notes: This figure shows the average age of non-steroid anti-cancer drugs that oncologists prescribe over time by level of specialization. Anti-cancer drugs are defined as unique World Health Organization Anatomical Therapeutic Class (ATC) codes that fall into the list of anti-cancer drugs from Pantziarka et al. (2021). Drug age is measured as the time since FDA approval of the ATC code. Drugs are weighted by the number of patients the oncologist prescribed the drug to. Oncologists are grouped into quantile bins based on the percentile of their HHI, where a higher HHI indicates a more specialized oncologist. Within each quantile bin, the sample is restricted to oncologists who prescribed anti-cancer drugs to at least 11 patients across a 20% sample of Part B and Part D data. HHI is calculated using shares of evaluation and management work Relative Value Units over types of cancer, and the percentile is calculated among the population of oncologists with at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims.

Figure A.8: Total dispensing fees for Part B physician administered drugs vary widely by cancer type



Notes: Figure A.8 shows the average Medicare Part B dispensing fees paid to oncologists per patient by cancer type in 2014. Dispensing fees are estimated using 6% of average sales price in the last quarter of 2014. The sample contains patients with a diagnosis of cancer in a 20% sample of 2014 Medicare Part B claims.

Table A.1: The length of cancer treatment guidelines over time

Year Site	2002	2009	2014	2020
Anus	9	22	32	55
Bladder	40	47	65	113
Brain and Spinal Cord	51	70	118	165
Breast	107	200	301	364
Cervix	24	31	67	86
Colorectal	38	145	217	355
Esophagus	19	48	125	154
Gastric and GE Junction	18	35	103	125
Head and Neck	86	116	182	233
Hodgkin's Lymphoma	30	39	74	85
Kidney	14	20	35	64
Liver and Biliary Tree	32	35	110	151
Lung and Bronchus	68	107	188	294
Lymphocytic Leukemia	0	0	96	141
Mast Cell	0	0	0	67
Mediastinum	0	0	30	41
Mesothelioma	0	0	30	46
Multiple Myeloma	27	42	66	95
Myelodysplastic and Myeloproliferative Disorders	22	32	53	192
Myeloid Leukemia	71	83	168	218
NHL	49	155	471	658
Neuroendocrine	46	64	98	141
Occult	30	43	55	65
Ovary and Adnexa	26	44	83	371
Pancreas	23	55	109	160
Placenta	0	0	0	40
Prostate	40	77	120	230
Sarcoma	56	92	196	281
Skin	58	82	144	465
Small Intestine	0	0	0	48
Testes	34	30	55	78
Thyroid	69	73	117	136
Uterus	42	45	81	108
All	**1,129**	**1,832**	**3,589**	**5,825**

Notes: Figure A.1 displays the page count of the first set of treatment guidelines in 2002, 2009, 2014, and 2020 from NCCN Guidelines. The NCCN considered the 2002 guidelines to be a "complete library" of oncology guidelines (NCCN, 2023*b*). However, in the years since, they have added guidelines for additional types of cancer.

Table A.2: For specialized oncologists, physician characteristics become substantially less predictive of exposure to knowledge growth after recentering

	(1) Unadjusted exposure to knowledge growth (Δk_i)	(2) Remove physician scope variation ($\Delta \tilde{k}_i$)	(3) Remove cancer prevalence variation ($\Delta \tilde{k}_i$)	(4) Remove endpoint type variation ($\Delta \tilde{k}_i$)
Market population (millions)	-29.35	-11.02	-5.809	-0.842
	(4.901)	(1.880)	(1.596)	(1.010)
Count oncologists at main org	-1.143	-0.262	0.121	0.109
	(0.288)	(0.110)	(0.0866)	(0.0494)
Physician age	0.305	0.836	0.171	-0.484
	(0.985)	(0.422)	(0.378)	(0.246)
Patient Count	5.168	0.720	-0.302	-0.129
	(0.352)	(0.120)	(0.115)	(0.0816)
Observations	988	988	988	988
R-squared	0.286	0.090	0.019	0.012
Recentered		Yes	Yes	Yes

Notes: This table shows the results of a regression of exposure to knowledge growth on observable physician characteristics for specialized oncologists. Specialized oncologists are defined as oncologists in the top 25% of the specialization distribution in 2009, measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. The physician characteristics include the population of the physicians' Hospital Referral Region in millions, the number of other oncologists at the Tax IDs that they work at, and physicians' patient volume in 2008. In Column 1, the outcome variable is raw exposure to knowledge growth, measured as the total increase in guideline page counts for cancer types from 2009 to 2014 for cancer types the physician managed in 2008. In Columns 2-4, the outcomes are increasingly recentered measures of exposure to knowledge growth. In recentering, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. In Column 2, predicted knowledge growth is the same for each cancer and estimated as mean knowledge growth per cancer type, weighted by the number of physicians who treat that cancer. In Column 3, the predicted knowledge growth for each type of cancer is estimated using on a Poisson regression of the number of Medicare patients with the cancer types in 2008. Column 4 is similar, but predicted knowledge growth also adds indicator for if the cancer type has well-established surrogate endpoints and the interaction between the number of Medicare patients and surrogate endpoints. The full recentering procedure removes variation in exposure to knowledge growth that arises become some physicians have a broader scope and because some cancer types predictably have more innovation. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014. All regression standard errors are robust.

Table A.3: For general oncologists, physician characteristics become substantially less predictive of exposure to knowledge growth after recentering

	(1) Unadjusted exposure to knowledge growth (Δk_i)	(2) Remove physician scope variation ($\Delta \tilde{k}_i$)	(3) Remove cancer prevalence variation ($\Delta \tilde{k}_i$)	(4) Remove endpoint type variation ($\Delta \tilde{k}_i$)
Market population (millions)	-3.707	-0.368	-0.0355	0.474
	(1.371)	(0.677)	(0.758)	(0.640)
Count oncologists at main org	0.00466	-0.0177	0.00903	0.000144
	(0.0559)	(0.0282)	(0.0337)	(0.0302)
Physician age	-0.843	0.171	0.183	-0.0190
	(0.282)	(0.139)	(0.161)	(0.138)
Patient Count	3.635	0.0565	-0.238	0.000888
	(0.0999)	(0.0405)	(0.0438)	(0.0392)
Observations	2,960	2,960	2,960	2,960
R-squared	0.409	0.002	0.009	0.000
Recentered		Yes	Yes	Yes

Notes: This table shows the results of a regression of exposure to knowledge growth on observable physician characteristics for general oncologists. General oncologists are defined as oncologists in the bottom 75% of the specialization distribution in 2009, measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. The physician characteristics include the population of the physicians' Hospital Referral Region in millions, the number of other oncologists at the Tax IDs they work at, and physicians' patient volume in 2008. In Column 1, the outcome variable is raw exposure to knowledge growth, measured as the total increase in guideline page counts for cancer types from 2009 to 2014 for cancer types the physician managed in 2008. In Columns 2-4, the outcomes are increasingly recentered measures of exposure to knowledge growth. In recentering, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. In Column 2, predicted knowledge growth is the same for each cancer and estimated as mean knowledge growth per cancer type, weighted by the number of physicians who treat that cancer. In Column 3, the predicted knowledge growth for each type of cancer is estimated using on a Poisson regression of the number of Medicare patients with the cancer types in 2008. Column 4 is similar, but predicted knowledge growth also adds indicator for if the cancer type has well-established surrogate endpoints and the interaction between the number of Medicare patients and surrogate endpoints. The full recentering procedure removes variation in exposure to knowledge growth that arises become some physicians have a broader scope and because some cancer types predictably have more innovation. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014. All regression standard errors are robust.

Table A.4: Effect of knowledge growth on specialization for all oncologists

	(1) Knowledge Growth Recentered ($\Delta\tilde{k}_i$)	(2) Restrict Sample	(3) Add Age FE	(4) Add Market FE	(5) Add Pat Count FE	(6) Add Org Size FE
Increase in Guideline Length (100 Pages)	73.74	71.95	71.62	87.93	88.79	89.36
	(31.27)	(31.43)	(31.41)	(32.72)	(32.71)	(32.67)
Observations	3,950	3,931	3,931	3,931	3,931	3,931
R-squared	0.002	0.002	0.002	0.087	0.087	0.088
Age FE			Yes	Yes	Yes	Yes
Market FE				Yes	Yes	Yes
Org Size FE					Yes	Yes
Pat Count FE						Yes

Notes: This table shows regression estimates for the relationship between exposure to knowledge growth and change in specialization for all oncologists. Specialization is measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. The measure of exposure to knowledge growth is recentered to control for variation in physician scope and field expected knowledge growth. Column 1 shows the coefficient estimate for the univariate regression. The regression standard errors are robust. Column 2 restricts the sample to physicians where we observe age, hospital referral region (HRR) and the number of other oncologists at their Tax IDs. Columns 3, 4, and 5 add fixed effects for the physicians age group, hospital referral region (HRR) and the number of other oncologists at their Tax IDs, respectively. In the recentered measure of exposure to knowledge growth, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. Raw exposure to knowledge growth is calculated as the total increase in the page count of NCCN Guidelines over 2009 to 2014 for the types of cancers a physician treats in 2008. Predicted exposure is estimated based on the cancer types a physician treats in 2008 and the predicted, rather than realized, growth in knowledge over the next five years. The predicted knowledge growth for each type of cancer is estimated based on a Poisson regression of the number of Medicare patients with the cancer types in 2008, an indicator for if the cancer type has well-established surrogate endpoints, and the interaction of the two. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014.

Table A.5: Effect of knowledge growth on number of unique cancer types managed for all oncologists

	(1) Knowledge Growth Recentered ($\Delta\tilde{k}_i$)	(2) Restrict Sample	(3) Add Age FE	(4) Add Market FE	(5) Add Pat Count FE	(6) Add Org Size FE
Increase in Guideline Length (100 Pages)	-0.141	-0.145	-0.156	-0.213	-0.233	-0.233
	(0.0800)	(0.0801)	(0.0798)	(0.0817)	(0.0812)	(0.0813)
Observations	3,950	3,931	3,931	3,931	3,931	3,931
R-squared	0.001	0.001	0.018	0.143	0.164	0.164
Age FE			Yes	Yes	Yes	Yes
Market FE				Yes	Yes	Yes
Pat Count FE					Yes	Yes
Org Size FE						Yes

Notes: This table shows regression estimates for the relationship between exposure to knowledge growth and the number of cancers managed for specialized oncologists. Specialized oncologists are defined as oncologists in the top 25% of the specialization distribution in 2009, measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. The measure of exposure to knowledge growth is recentered to control for variation in physician scope and field expected knowledge growth. Column 1 shows the coefficient estimate for the univariate regression. The regression standard errors are robust. Column 2 restricts the sample to physicians where we observe age, hospital referral region (HRR) and the number of other oncologists at the Tax IDs. Columns 3, 4, and 5 add fixed effects for the physicians age group, HRR and the number of other oncologists at their Tax IDs, respectively. In the recentered measure of exposure to knowledge growth, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. Raw exposure to knowledge growth is calculated as the total increase in the page count of NCCN Guidelines over 2009 to 2014 for the types of cancers a physician treats in 2008. Predicted exposure is estimated based on the cancer types a physician treats in 2008 and the predicted, rather than realized, growth in knowledge over the next five years. The predicted knowledge growth for each type of cancer is estimated based on a Poisson regression of the number of Medicare patients with the cancer types in 2008, an indicator for if the cancer type has well-established surrogate endpoints, and the interaction of the two. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014.

Table A.6: Effect of knowledge growth on specialization for ex-ante specialized oncologists (estimated with 25 patients per oncologist)

	(1)	(2)	(3)
Increase in Guideline Length (100 Pages)	131.5	169.2	-47.37
	(78.55)	(93.57)	(103.2)
Observations	988	988	988
R-squared	0.003	0.003	0.000
Knowledge Exposure Sample	25 pat	All pat	25 pat
HHI Sample	All pat	25 pat	25 pat

Notes: This table shows regression estimates for the relationship between exposure to knowledge growth and change in specialization for oncologists who start in the top 25% of the initial specialization distribution. Specialization is measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. The measure of exposure to knowledge growth is recentered to control for variation in physician scope and field expected knowledge growth. Column 1 shows the regression with exposure to knowledge growth calculated using only 25 random patients. Column 2 shows the regression with $H\hat{H}I_i$ calculated using only 25 random patients. Column 3 shows the regression with both values calculated using only 25 random patients In the recentered measure of exposure to knowledge growth, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. Raw exposure to knowledge growth is calculated as the total increase in the page count of NCCN Guidelines over 2009 to 2014 for the types of cancers a physician treats in 2008. Predicted exposure is estimated based on the cancer types a physician treats in 2008 and the predicted, rather than realized, growth in knowledge over the next five years. The predicted knowledge growth for each type of cancer is estimated based on a Poisson regression of the number of Medicare patients with the cancer types in 2008, an indicator for if the cancer type has well-established surrogate endpoints, and the interaction of the two. Unless otherwise specified, all physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014.

Table A.7: Effect of knowledge growth on specialization for ex-ante general oncologists (estimated with 25 patients per oncologist)

	(1)	(2)	(3)
Increase in Guideline Length (100 Pages)	29.90	103.2	128.1
	(29.44)	(47.18)	(47.94)
Observations	2,962	2,962	2,962
R-squared	0.000	0.002	0.003
Knowledge Exposure Sample	25 pat	All pat	25 pat
HHI Sample	All pat	25 pat	25 pat

Notes: This table shows regression estimates for the relationship between exposure to knowledge growth and change in specialization for oncologists who start in the lower 75% of the initial specialization distribution. Specialization is measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. The measure of exposure to knowledge growth is recentered to control for variation in physician scope and field expected knowledge growth. Column 1 shows the regression with exposure to knowledge growth calculated using only 25 random patients. Column 2 shows the regression with $H\hat{H}I_i$ calculated using only 25 random patients. Column 3 shows the regression with both values calculated using only 25 random patients In the recentered measure of exposure to knowledge growth, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. Raw exposure to knowledge growth is calculated as the total increase in the page count of NCCN Guidelines over 2009 to 2014 for the types of cancers a physician treats in 2008. Predicted exposure is estimated based on the cancer types a physician treats in 2008 and the predicted, rather than realized, growth in knowledge over the next five years. The predicted knowledge growth for each type of cancer is estimated based on a Poisson regression of the number of Medicare patients with the cancer types in 2008, an indicator for if the cancer type has well-established surrogate endpoints, and the interaction of the two. Unless otherwise specified, all physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014.

Table A.8: Effect of knowledge growth on specialization for all oncologists (estimated with 25 patients per oncologist)

	(1)	(2)	(3)
Increase in Guideline Length (100 Pages)	43.15	114.0	86.88
	(29.24)	(42.31)	(43.97)
Observations	3,950	3,950	3,950
R-squared	0.001	0.002	0.001
Knowledge Exposure Sample	25 pat	All pat	25 pat
HHI Sample	All pat	25 pat	25 pat

Notes: This table shows regression estimates for the relationship between exposure to knowledge growth and change in specialization for all oncologists. Specialization is measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. The measure of exposure to knowledge growth is recentered to control for variation in physician scope and field expected knowledge growth. Column 1 shows the regression with exposure to knowledge growth calculated using only 25 random patients. Column 2 shows the regression with $H\hat{H}I_i$ calculated using only 25 random patients. Column 3 shows the regression with both values calculated using only 25 random patients In the recentered measure of exposure to knowledge growth, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. Raw exposure to knowledge growth is calculated as the total increase in the page count of NCCN Guidelines over 2009 to 2014 for the types of cancers a physician treats in 2008. Predicted exposure is estimated based on the cancer types a physician treats in 2008 and the predicted, rather than realized, growth in knowledge over the next five years. The predicted knowledge growth for each type of cancer is estimated based on a Poisson regression of the number of Medicare patients with the cancer types in 2008, an indicator for if the cancer type has well-established surrogate endpoints, and the interaction of the two. Unless otherwise specified, all physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014.

Table A.9: Effect of knowledge growth, estimated using count of guideline references, on specialization for ex-ante specialized oncologists

	(1) Knowledge Growth Recentered ($\Delta\tilde{k}_i$)	(2) Restrict Sample	(3) Add Age FE	(4) Add Market FE	(5) Add Pat Count FE	(6) Add Org Size FE
Increase in Guideline Citations (100 Citations)	64.91	69.41	70.88	84.97	85.27	83.76
	(18.49)	(18.83)	(18.89)	(21.37)	(21.45)	(21.61)
Observations	988	924	924	924	924	924
R-squared	0.014	0.016	0.017	0.182	0.187	0.188
Age FE			Yes	Yes	Yes	Yes
Market FE				Yes	Yes	Yes
Org Size FE					Yes	Yes
Pat Count FE						Yes

Notes: This table shows regression estimates for the relationship between exposure to knowledge growth and change in specialization for oncologists who start in the top 25% of the initial specialization distribution. Specialization is measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. The measure of exposure to knowledge growth is recentered to control for variation in physician scope and field expected knowledge growth. Column 1 shows the coefficient estimate for the univariate regression. The regression standard errors are robust. Column 2 restricts the sample to physicians where we observe age, hospital referral region (HRR) and the number of other oncologists at their Tax IDs. Columns 3, 4, and 5 add fixed effects for the physicians age group, hospital referral region (HRR) and the number of other oncologists at their Tax IDs, respectively. In the recentered measure of exposure to knowledge growth, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. Raw exposure to knowledge growth is calculated as the total increase in the count of guideline citations over 2009 to 2014 for the types of cancers a physician treats in 2008. Predicted exposure is estimated based on the cancer types a physician treats in 2008 and the predicted, rather than realized, growth in knowledge over the next five years. The predicted knowledge growth for each type of cancer is estimated based on a Poisson regression of the number of Medicare patients with the cancer types in 2008, an indicator for if the cancer type has well-established surrogate endpoints, and the interaction of the two. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014.

Table A.10: Effect of knowledge growth, estimated using count of guideline references, on specialization for ex-ante general oncologists

	(1) Knowledge Growth Recentered ($\Delta\tilde{k}_i$)	(2) Restrict Sample	(3) Add Age FE	(4) Add Market FE	(5) Add Pat Count FE	(6) Add Org Size FE
Increase in Guideline Citations (100 Citations)	15.62	16.58	16.60	19.26	18.94	18.64
	(7.250)	(7.263)	(7.268)	(7.785)	(7.770)	(7.742)
Observations	2,962	2,937	2,937	2,937	2,937	2,937
R-squared	0.002	0.002	0.002	0.108	0.108	0.110
Age FE			Yes	Yes	Yes	Yes
Market FE				Yes	Yes	Yes
Org Size FE					Yes	Yes
Pat Count FE						Yes

Notes: This table shows regression estimates for the relationship between exposure to knowledge growth and change in specialization for oncologists who start in the bottom 75% of the initial specialization distribution. Specialization is measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. The measure of exposure to knowledge growth is recentered to control for variation in physician scope and field expected knowledge growth. Column 1 shows the coefficient estimate for the univariate regression. The regression standard errors are robust. Column 2 restricts the sample to physicians where we observe age, hospital referral region (HRR) and the number of other oncologists at their Tax IDs. Columns 3, 4, and 5 add fixed effects for the physicians age group, hospital referral region (HRR) and the number of other oncologists at their Tax IDs, respectively. In the recentered measure of exposure to knowledge growth, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. Raw exposure to knowledge growth is calculated as the total increase in the count of guideline citations over 2009 to 2014 for the types of cancers a physician treats in 2008. Predicted exposure is estimated based on the cancer types a physician treats in 2008 and the predicted, rather than realized, growth in knowledge over the next five years. The predicted knowledge growth for each type of cancer is estimated based on a Poisson regression of the number of Medicare patients with the cancer types in 2008, an indicator for if the cancer type has well-established surrogate endpoints, and the interaction of the two. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014.

Table A.11: Effect of knowledge growth, estimated using count of guideline references, on specialization for all oncologists

	(1) Knowledge Growth Recentered ($\Delta\tilde{k}_i$)	(2) Restrict Sample	(3) Add Age FE	(4) Add Market FE	(5) Add Pat Count FE	(6) Add Org Size FE
Increase in Guideline Citations (100 Citations)	26.46	26.28	26.23	29.94	30.28	30.25
	(7.281)	(7.312)	(7.317)	(7.720)	(7.728)	(7.721)
Observations	3,950	3,931	3,931	3,931	3,931	3,931
R-squared	0.004	0.004	0.004	0.089	0.089	0.090
Age FE			Yes	Yes	Yes	Yes
Market FE				Yes	Yes	Yes
Org Size FE					Yes	Yes
Pat Count FE						Yes

Notes: This table shows regression estimates for the relationship between exposure to knowledge growth and change in specialization for all oncologists. Specialization is measured using a Herfindahl—-Hirschman Index (HHI) over types of cancer managed. The measure of exposure to knowledge growth is recentered to control for variation in physician scope and field expected knowledge growth. Column 1 shows the coefficient estimate for the univariate regression. The regression standard errors are robust. Column 2 restricts the sample to physicians where we observe age, hospital referral region (HRR) and the number of other oncologists at their Tax IDs. Columns 3, 4, and 5 add fixed effects for the physicians age group, hospital referral region (HRR) and the number of other oncologists at their Tax IDs, respectively. In the recentered measure of exposure to knowledge growth, predicted exposure to knowledge growth is subtracted from raw exposure to knowledge growth. Raw exposure to knowledge growth is calculated as the total increase in the count of guideline citations over 2009 to 2014 for the types of cancers a physician treats in 2008. Predicted exposure is estimated based on the cancer types a physician treats in 2008 and the predicted, rather than realized, growth in knowledge over the next five years. The predicted knowledge growth for each type of cancer is estimated based on a Poisson regression of the number of Medicare patients with the cancer types in 2008, an indicator for if the cancer type has well-established surrogate endpoints, and the interaction of the two. All physicians in the sample are hematologist-oncologists, oncologists, or hematologists who saw at least 25 unique cancer patients in a 20% sample of Medicare Part B outpatient evaluation and management claims in 2008, 2009 and 2014.

B Data Appendix

This appendix provides additional detail and discussion about data construction. Section B.2 provides details and examples about the construction of the custom groupings for cancer diagnosis codes. Section B.3 provides details and examples about the construction of the indicator for academic status for physicians.

B.1 Cancer types with no guidelines

One limitation of the data is that some rare cancers lack guidelines in early years despite NCCN considering their 2002 guidelines to be a "complete library" of clinical guidelines. This limitation is evident Appendix Table A.1, which shows that a handful of cancer types lack guidelines in 2002, 2009, and 2014 One example is mesothelioma, a rare type of cancer that did not have guidelines until 2010. A 2010 news article about the release of the first mesothelioma guidelines notes that the guideline authors' found the guidelines "challenging to prepare" in part because of limited "data available" about how to treat the condition (Mulcahy, 2010). This example illustrates a key aspect of the guideline measure: zero pages of guidelines largely reflect that there is near-zero evidence about how to treat that specific type of cancer. However, patients with mesothelioma were being treated in some fashion prior to 2010, likely based on general principles of cancer treatment. Fortunately, the empirical strategy does not require the assumption that zero pages of guidelines reflect zero knowledge, only that more pages reflect more knowledge. However, the introduction of new guidelines may create a jump the number of pages of guidelines that overstates the relative knowledge growth in these cancers. This measurement challenge impacts a very small share of patients in the sample (2.80% in 2009).

B.2 Cancer Groupings

I create a set of custom diagnosis code groupings that align with how cancers are grouped in NCCN Guidelines. This section discusses why and how I create these custom groupings. The code and mappings to create these groupings can be freely shared upon request.

I created this grouping because ICD-9 diagnosis codes are too granular to be used directly as types of cancer treated. For example, they contain separate codes for breast cancer in the right breast versus the left breast. I wish to group these into the category of "breast cancer" to align with the NCCN Guidelines for breast cancer treatment.

To create the grouping, I start with a list of cancer diagnosis codes from the Clinical Classification Software (CCS). The CCS also has groupings to the level of cancer type, but these groupings have several challenges that I rectify with the custom grouping. First, I drop the CCS diagnosis codes for cancer treatment and consider only those for specific types of cancer. Second, I make the groupings more specific. For example, the CCS groupings classify colon cancer of "unknown behavior" as an "unknown" cancer type. I re-classify this as colon cancer. Third, I make the custom grouping more granular to align with the NCCN Guidelines. For example, I divide the CCS category of "Leukemia" into "Myeloid Leukemia" and "Lymphocytic Leukemia" (NCCN Guidelines are even further split by acute and chronic within each subtype of leukemia). Lastly, I rearrange groupings that are misaligned with clinical groupings. For example, the CCS groupings group rectal cancer with anal cancer. However, NCCN Guidelines group rectal cancer with colon cancer. The custom grouping follows the NCCN Guidelines' approach, with a single category for colon and rectal cancer.

Notably, cancer groupings in NCCN Guidelines have become more granular over time. I group NCCN Guidelines together to form cancer type groups, defaulting to the broadest (and typically oldest) grouping seen in NCCN Guidelines. For example, the NCCN Guidelines contain a single set of guidelines for "Non-Hodgkins'Lymphoma" through 2016. After 2016, these guidelines are divided into separate guidelines for B-cell lymphomas, T-cell lymphomas, Lymphoplasmacytic Lymphomas, and others. I group all of these into single category of

Non-Hodgkins'Lymphoma in order to have a stable grouping over time.

I also group screening and diagnosis NCCN Guidelines with treatment NCCN Guidelines where relevant when counting guidelines lengths and citation counts. For most cancers, all types of recommendations are in a single guideline. However, breast and prostate cancer have separate NCCN Guidelines for screening and diagnosis than for treatment. I group all types of NCCN Guidelines for breast cancer and all types of NCCN Guidelines for prostate cancer together. This grouping ensure breast and prostate cancer measures of length are analogous to the measures for other types of cancers.

These mappings are based on ICD-9 codes. In mid-2015, the claims data switches from ICD-9 to ICD-10 codes. From 2015 onwards, I map from ICD-10 to ICD-9 codes, then to the custom cancer grouping.

Notably, the ICD-based groupings are based on the anatomical origin of the tumor, which mirrors the first-order dimension of cancer subtyping. For example, breast cancers are cancers where the origin tumor site is the breast, regardless of where the cancer has spread. Within each anatomical origin, there are further subclassifications along numerous other dimensions, e.g., stage and biomarkers. I cannot observe these other dimensions in the diagnosis codes. However, I do not need to observe them as NCCN Guidelines are also organized based on the anatomical origin of tumors.

In addition, I am reassured that anatomic origin is the appropriate high-level grouping because NCCN Guidelines are also organized anatomically. This observation suggests that the grouping of cancer types aligns with how the field of oncology initially categorizes cancer (there are many dimensions of additional subtyping by size, spread, biomarkers, and other features). This anatomical approach also aligns with how sub-specialized oncologists at the University of Chicago describe their focus area, i.e., using anatomic designations such as “breast cancer oncologist”, “GI oncologists”, or “thoracic oncologist”.

B.3 Labeling Academic Physicians

This section discusses how I labeled physicians as academic or non-academic. The resulting measure of academic status is a broadly inclusive proxy of academic status rather than a definitive measure. These limitations arise from the inherent challenges of binarizing the diversity of organizational arrangements in oncology and the limitations of the data. I detail the process below, which broadly emulates the approach from Welch and Bindman (2016).

I use the Medicare Data on Provider Practice and Specialty (MD-PPAS) from 2008 to proxy academic status. The data contains the name associated with the tax ID number from which the physician billed the most services in 2008. I assign academic status based on the name of this tax ID number.

A tax ID name must meet at least one of several possible criteria to be categorized as academic. The tax ID is considered academic if it contains the name of a recognizable academic institution or other words that signal academic status (e.g. "university", "college"). I also consider a tax ID to be academic if its name matches the state and name of record for an academic institution's physician group. I also consider a tax ID to be academic if the institution named was affiliated with an academic center or university on its website as of May 2024, the time when this determination of academic status was made. I use affiliation in May 2024 rather than 2008 as I have more limited tools to determine past affiliation status. Finally, I also consider a tax ID number to be academic if the institution named has a Hematology-Oncology fellowship training program, even if it is not affiliated with a university.

The resulting measure of academic status is broadly inclusive. For example, satellite locations of the Mayo Clinic in Florida and Arizona are counted as academic, as are regional affiliates of Memorial Sloan Kettering Cancer Center. Similarly, the Northshore University Health System is labeled as academic since its name contains the phrase "University" and is currently affiliated with the University of Chicago. This measure of academic status prioritizes sensitivity (including all plausibly academic physicians) over specificity (definitively

excluding all non-academic physicians from the label of "academic").

C Measurement Error in Physician Specialization

This section discusses measurement error in the main outcome variable, physician specialization. Estimated physician specialization has measurement error with a positive expected value. Fortunately, the recentering procedure used in the main estimates removes the correlation of this measurement error with the treatment variable (exposure to knowledge growth). In addition, this measurement error does not change the paper's key findings.

C.1 Measurement Error in Exposure to Knowledge Growth

Recall that I do not observe all patients nationally but rather a 20% sample of patients enrolled in Medicare Part B. Not observing all patients may lead to measurement error in exposure to knowledge growth, both unadjusted (Δk_i) and recentered ($\Delta \tilde{k}_i$). A key concern is that the degree of finite-sample measurement error is correlated with a physician's patient volume (i.e. how finite the sample is). Fortunately, patient volume is balanced across the recentered measure of exposure to knowledge growth, $\Delta \tilde{k}_i$ (recall Table 5), which implies that the degree of finite-sample measurement error is uncorrelated with the main treatment variable.

C.2 Measurement Error in Physician Specialization

Measurement error occurs because I observe only a subset of the patients a physician treats in a year. Recall that physician specialization is measured by a Herfindahl–Hirschman Index (HHI) over the shares of time that an oncologist spends managing each type of cancer: $HHI_i = 10,000 \sum_C \left(\frac{T_{ic}}{T_i}\right)^2$. The estimated value of HHI_i will be biased upwards on average in a finite subsample of patients. To illustrate, suppose physician i has $HHI_i = 2500$ when calculated over all patients. Suppose I observe only one patient; mechanically, I will estimate that $\hat{HHI}_i = 10000$. Suppose, I observe only two patients; mechanically, the minimum HHI I could estimate is $\hat{HHI}_i = 5000$. As the number of patients increases, the expected value of

$H\hat{H}I_i$ approaches HHI_i from above. Recall that I observe only a 20% sample of Medicare patients treated by a physician rather than all patients. As such, I estimate HHI_i with finite sample measurement error. This measurement error will be larger for low-volume physicians and decline as patient volume increases.

Figure A.6 demonstrates that $H\hat{H}I_i$ declines with volume and that some of the decline is attributable to decreases in measurement error. I plot mean $H\hat{H}I_i$ by patient volume in 1999, 2009, and 2019. $H\hat{H}I_i$ declines with volume in all three years. How much of this decline comes from reductions in measurement error, and how much reflects real specialization among low-volume physicians? For example, the observed relationship could reflect real specialization differences if low-volume physicians tend to be academic sub-specialists who spend half their time treating one type of cancer and half their time researching that cancer.

To distinguish between measurement error and real differences in specialization, I also plot $H\hat{H}I_i$ estimated with 25 random patients per doctor. Using 25 random patients inflates the finite sample measurement error for physicians who see more than 25 patients a year so that $H\hat{H}I_i$ is measured with an equal degree of error regardless of observed physician volume.

There are several key findings. First, small sample noise is not mean zero. The HHI calculated with 25 patients is consistently higher than the HHI calculated with all patients. Second, the relationship between volume and HHI changes over time. In 1999, low-volume and high-volume doctors have similar HHIs once I measure HHI with 25 patients for everyone. This pattern suggests that the decline in HHI with volume is largely a product of volume-dependent measurement error. However, in 2009 and 2019, HHI does decline with volume, even after equalizing volume-dependent measurement error. This pattern implies that specialists are genuinely lower-volume physicians in more recent years.

Next, I consider the impact of volume-dependent measurement error on my main identification strategy. A key concern for estimation is that patient volume, and therefore measurement error in specialization, is correlated with the treatment variable, creating omitted

variables bias in the main regression estimates. Table 5 shows that the raw measure of exposure to knowledge growth is correlated with patient volume and, therefore, likely correlated with measurement error. Fortunately, the recentering procedure removes the correlation between exposure to knowledge growth and physician volume and, therefore, measurement error. If so, non-zero measurement error should not contribute to omitted variables bias in the main specifications. In addition, the main regression results are robust to adding fixed effects for patient volume (Tables 6 and 7), suggesting that patient volume, and therefore volume-dependent measurement error, is not a source of omitted variables bias.

C.3 Effect of equalizing volume-dependent measurement error in main results

Next, I investigate the impact of volume-dependent measurement error on the main findings. To do so, I replicate the key descriptive and regression results with a measure of specialization ($\hat{HHI}_i$) and recentered exposure to knowledge growth ($\Delta\tilde{k}_i$) calculated over 25 random patients per physician instead of all observed patients. This approach inflates measurement error for high-volume physicians and thus equalizes measurement error across low- and high-volume physicians.

First, I consider the impact on the observed trends over time in specialization. Figure A.3 plots the distribution of HHI over time, similar to figure 3, but calculates HHI using only 25 random patients per physician rather than all patients. The overall patterns are essentially unchanged. The 10th through 75th percentiles of specializations are largely flat or declining. The 90th percentile of specialization increases substantially over time, consistent with rising top specialization.

Finally, I replicate the main regressions using only 25 patients per physician. Table A.6 shows the regression of exposure to knowledge growth on specialization for ex-ante specialized physicians. Column 1 shows the regression with exposure to knowledge growth calculated using only 25 random patients. Column 2 shows the regression with $\hat{HHI}_i$ calculated using

only 25 random patients. Column 3 shows the regression with both exposure to knowledge growth and $H\hat{H}I_i$ calculated using only 25 random patients. The estimated effects are similar to the main results, with the same direction and positive magnitude. Table A.7 shows the same results for ex-ante generalist physicians. The estimates are all indistinguishable from zero, as in the main results. The similar estimates suggest that measurement error does not substantively impact or change the central findings.