

Labor Markets and Technological Change: Evidence from Electronic Health Records

January 5, 2026

Abstract

Technology adoption in health care has the potential to improve patient outcomes, and in some cases, decrease costs. However, new technologies can also be burdensome for physicians. I add to our understanding of how new technology can affect its users by investigating changes in physician behavior as a result of a major technology shift in US hospitals, electronic health records (EHRs). I treat EHR adoption in hospitals as an exogenous shock to hospitalists, and estimate average group time treatment effects on various labor market outcomes. I find that hospitalists are more likely to leave clinical work altogether, and more likely to shift work to an office setting due to EHRs. For hospitalists that continue working in the hospital, EHRs lead to an increase in patients seen. These findings inform the full costs and benefits of technology adoption, relevant to both technology adopters and policymakers who incentivize technology adoption.

Keywords: Electronic Health Record, Physician Behavior, Technology Adoption

JEL Codes: I11, J24, O33

1 Introduction

For decades technological innovation has altered production through changes in labor demand, wages, and other factors. While modern technology advances rapidly, it is unclear how policy-makers should engage with its implementation. This is particularly important in health care, where new technologies also impact patient outcomes. A recent technology roll-out that altered the US health care system dramatically was the electronic health record (EHR). Medical entities nationwide implemented EHRs from 2008–2015, with basic EHR capability rising from 9 to 84 percent of hospitals. In this paper, I investigate the effect of EHR adoption in hospitals on several behaviors of physicians working in hospitals, hereafter referred to as hospitalists.

EHRs are computer systems that house digital medical records, and they also have the capability to be used for more complex functionalities. EHRs became increasingly relevant after 2009, when the HITECH Act subsidized hospitals and practices to “meaningfully” use EHRs (Civil Rights, 2017).¹ President Obama stated in 2009, “To improve the quality of our health care while lowering its cost, we will make the immediate investments necessary to ensure that, within five years, all of America’s medical records are computerized.” (Chang, Childs, and Grayson, 2009). The push for EHR use stemmed from a widespread expectation that EHRs would improve quality of health care while decreasing costs, the gold standard in health policy. For example, a 2005 study estimated hundreds of billions of dollars saved if health information technology were to be fully implemented (Hillestad et al., 2005).

Physicians, as the primary users of EHRs, play an important role in whether the expected benefits come to fruition. Daily tasks change drastically with a new system; some physicians report lower job satisfaction and higher stress levels. Senior physicians in particular “loathe the cumbersome, time-consuming data entry that comes with using EHR.” (Collier, 2017a). When looking solely at hospitalists, 63% of survey respondents found EHRs time consuming and inefficient (Czernik et al., 2022). The frustration of a new technology raises the cost of working in settings that use it, which may lead physicians on the margin to make behavioral changes such as exiting the labor market altogether or shifting towards alternative work settings. The extent to which this frustration imposes a meaningful cost is unknown, making physician response an empirical question. For a hospitalist, the hospital’s decision to adopt an EHR is plausibly exogenous. Thus, I use a difference-in-differences research design to estimate the effect of EHR adoption on hospitalist labor market choices: (1) exiting clinical work altogether, measured based on no future claims,

¹Meaningful use (MU) standards proceeded in three stages over time. In Stage 1 (2010), MU focused on data capturing and sharing. In Stage 2, which began in late 2012, MU extended to using EHRs for patient incorporation and using the technology as a helper in care. Stage 3 went from 2014-2016 and focused on making data accessible across hospitals (*Understanding Meaningful Use*)

(2) substituting inpatient work to offices, measured by the location of claims, and (3) care/billing patterns, measured by the number of patients seen and claims filed per patient.

While hospitalists are an ideal setting to answer this question econometrically, they are also important in their own right. Approximately 44,000 physicians are employed as hospitalists in the US, and studies estimate that 75% of hospitals employ hospitalists (BLS, 2015; eCareers, 2022). Further, one third of Medicare admissions are seen by a hospitalist (Messler and Whitcomb, 2015). I construct a panel of hospitalists spanning from 2009-2017. Using Centers for Medicare and Medicaid Services (CMS) Shared Patient Data, I identify the hospital affiliation of each hospitalist. Then, I use the American Hospital Association (AHA) Survey and Health Information and Management Systems Society (HIMSS) to determine EHR adoption status, and thus connect hospitalists to EHR exposure. Finally, I use Medicare Data on Provider Practice and Specialty (MDPPAS) for information on the patients seen by each hospitalist. The independent variable of interest is binary, capturing exposure to an EHR.² Due to staggered treatment timing, I estimate group-time treatment effects of EHR exposure on the physician decisions defined above (Callaway and Sant’Anna, 2021).

This paper expands our understanding of the effects of health information technology. Despite a large number of case studies that find generally positive effects, i.e., improved patient outcomes and decreased cost (surveyed in Buntin et al., 2011), quantitative studies have found a mixture of results: outcomes for median patients do not change as a result of EHRs (Agha, 2014, McCullough, Parente, and Town, 2016, Meyerhoefer et al., 2016) newborns and severe patients experience improvement in health outcomes (A. R. Miller and Tucker, 2009, Freedman, Lin, and Prince, 2015, McCullough, Parente, and Town, 2016), and hospital costs only decrease 6 years after implementation, if at all (Agha, 2014, Dranove, Forman, et al., 2014). Relating EHRs to productivity, one study finds that nursing home productivity increases after adopting health IT (Hitt and Tambe, 2016), but another finds that physician productivity decreases by 11 percent due to EHRs being adopted in primary-care sites (Meyerhoefer et al., 2016). There is still much to be learned about EHRs effects on health care, specifically related to labor supply (Bronsolero, Doyle, and Van Reenen, 2021). I contribute to this literature first by considering individual physician labor market decisions as an outcome, which, to my knowledge, has not yet been examined. Further, I consider the primary time period in which EHRs were rapidly implemented, in part due to the HITECH Act (Dranove, Garthwaite, et al., 2015). Similar to past studies, I treat EHR implementation as a treatment variable in a difference-in-differences framework, improving on past strategies by estimating group-time effects across multiple years, avoiding problems in typical two-way fixed

²While there are ample nuances to EHRs, I abstract away from these as I focus on the initial shock of a new technology.

effects models with staggered treatment.

I find that EHR implementation led to a 20% increase in the likelihood of exiting clinical settings for hospitalists over age 59. Physicians less than 60 were 7% more likely to exit, possibly pointing to career switching. Limiting the sample to those who remained in clinical settings, I find that hospitalists were 30% more likely to start seeing patients in an office setting at the time of exposure. Finally, for hospitalists who do not change work setting, I find that EHRs lead to a patient count increase of at least 20 patients per year. This increase is large enough that it is not entirely driven by a redistribution of patients from those who left clinical settings. Underlying this analysis, I assume the decision for hospitals to implement EHRs is exogenous to individual hospitalists, and uncorrelated with other major changes to the hospitalist work environment. I mitigate concerns of the latter by providing evidence that no other observable hospital changes are correlated with EHR adoption on average. Additionally, the main findings are not sensitive to a plethora of robustness checks.

The remainder of the paper is organized as follows. Section 2 provides background on EHR adoption in the US and hospitalists. Section 3 outlines the data used and provides summary statistics. Section 4 details the empirical strategy and identification assumptions. Finally, Section 5 contains the results, and Section 6 rules out several alternative explanations.

2 Background

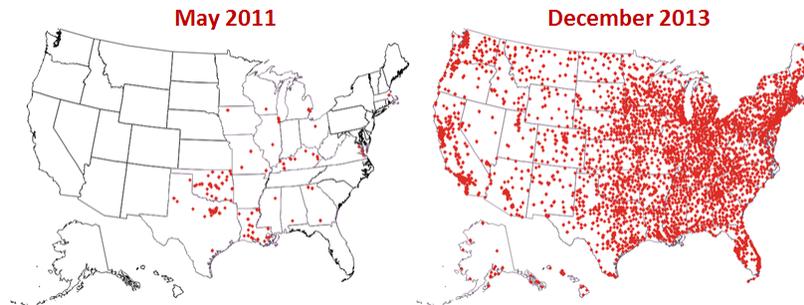
2.1 EHRs and Meaningful Use

EHRs have been an important feature of health care since the 1980s. Early in their existence, health care professionals perceived the technology as a complement to paper records, primarily deployed by large academic medical centers to improve billing and scheduling efficiency. Physicians did not interact directly with these early-generation EHRs and thus were not drastically affected by their implementation. As innovation made computers more portable, EHR usability increased, creating what is known as the “physician workspace”: a computer station for physicians to interface directly with an EHR to record patient updates. Despite usability improvements, physicians continued to view EHRs as purely complementary to paper due to burdensome data entry (Evans, 2016).

The HITECH Act, passed in 2009, designated \$27 billion in subsidies to entities using EHRs according to specific guidelines. These included having at least 80% of patients in the system, recording answers to specific questions, and ensuring privacy. Subsidies were allocated in stages: stage 1 (2011) focused on data collection; stage 2 (2012) extended to care support; and stage 3 (2014) extended to data sharing between practices. Hospitals without EHRs by 2015 faced penalties through lower Medicare reimbursements. Vast EHR adoption coincided with the rollout

of these subsidies, shown in a 75 percentage point increase in the number of hospitals with EHRs from 2009 to 2015 (Dashboard, 2019). However, subsidies did not necessarily cause all adoption. Dranove, Garthwaite, et al. (2015) estimate adoption would have been ten percentage points lower without the policy. Figure 1 shows a geographical comparison of US hospitals receiving stage 1 meaningful use subsidies in 2011 vs. 2013, revealing the nationwide expansion of technology.

Figure 1: Hospitals Receiving Meaningful Use Stage 1 Subsidy



Source: HealthIT.gov

On a daily basis, physicians spend approximately 23.7%, 17%, and 15.5% of their time on documentation, chart review, and inbox management, respectively—the most time-consuming aspects of EHR use (Arndt et al., 2017). If a physician works for 12 hours, they spend roughly 3.2 hours with patients and 5.9 hours interfacing with an EHR. Despite the time spent, physicians continue to report frustration over EHR use. The most time-consuming functions are also reported as the most frustrating (Dymek et al., 2021). Another common issue is the lack of usability, with functions taking minutes to locate or load. Hospitals often provide training upon adoption, but rapid upgrades and changes make familiarity difficult. While these findings focus on outpatient settings, the limited research on hospitalists shows similar patterns (Tipping et al., 2010). Together, these findings imply a meaningful additional cost imposed on physicians when new technology is introduced.

2.2 Hospitalists and Labor Market Decisions

This section outlines the potential responses of hospitalists to EHRs based on their incentives and institutions. The behaviors considered are closely related and should not be taken as independent. For example, hospitalists' choice of work setting depends on their decision to remain in clinical practice, and billing activity depends on work setting. Thus, I discuss the behaviors in stages of decision making.

For ease of notation, I define a hospitalist as any general practice physician who sees at least

2.2.1 Exiting Clinical Work

Physicians can stop seeing patients either by retiring or switching to non-clinical careers. While retirement decisions are typically planned and financially driven (Gustman and Steinmeier, 1986; Stock, Wise, et al., 1990), physicians often delay retirement for altruistic reasons, such as maintaining patient relationships (Collier, 2017b). Thus, the option to retire may exist for years before a physician actually exits clinical practice. Age plays a central role, with older physicians more likely to be influenced by technology shocks like EHR adoption. If EHRs reduce job satisfaction, they may push physicians, especially those near retirement, to leave clinical work. Alternatively, if EHRs improve daily tasks, they may encourage continued practice. Media accounts suggest that EHR implementation did in fact prompt some older physicians to retire (Ringel, 2019; Loria, 2020), though this remains empirically untested.

Physicians of any age may also leave clinical roles for careers in administration, teaching, consulting, or hospital management. While early career physicians are unlikely to fully retire, they may still leave clinical work if EHRs reduce job satisfaction. In a 2016 survey, 13

2.2.2 Work Setting

For physicians who remain practicing, different practice settings may be in the choice set. Specifically, if EHRs affect job satisfaction in hospitals, a hospitalist may shift more working hours toward or away from the hospital. This is particularly salient for hospitalists who contract time, as not all have binding contracts. A 2013 survey estimated that 36% of hospitalists also worked outside the hospital in primary care, long-term care, or hospice settings (Hospitalist, 2016). Many internists also split time across inpatient and office settings (Gray et al., 2022).

Even salaried hospitalists may shift to outpatient settings if their hospital owns such facilities. While office and inpatient services are not fully interchangeable, changes in claims from hospital to office settings can signal shifts in practice location. If EHRs influence how hospitals allocate physicians, this shift may reflect hospital-level decisions rather than physician preferences. My sample includes both salaried hospitalists and internists who split time across settings. Although I cannot observe individual physician contracts, I use hospital-level integration to explore whether changes reflect hospital or physician-driven decisions.

2.2.3 Care/Billing Activity

For those who do not exit or change work settings, it is natural to assess whether EHRs impacted patient care patterns. A main goal of the HITECH Act was to improve efficiency of care by reducing administrative burdens. However, EHRs could add daily tasks that hinder productivity, enhance efficiency, or both. I cannot directly analyze productivity without granular information

on patient sharing and time use, but I do observe total patient activity, an aspect of productivity. Whether changes to patient counts affect patient welfare is outside of the scope of this paper.

EHRs may influence hospitalist billing behavior in complex ways. Physicians often over-treat due to defensive medicine, patient demand, or limited access to records (Lyu et al., 2017). EHRs could reduce over-treatment by improving information access, decreasing repeated tests—and claims per patient. However, hospitals may also use EHRs to strategically increase billing.

3 Data

I construct a balanced hospitalist-level panel data set spanning from 2009-2017 which measures physician exposure to EHRs over time, several labor market outcomes, and other characteristics. I describe the different data sets used to construct the panel below, and I include a detailed outline of data and variable construction in Appendix A.

3.1 Sample Construction

I first gather physician information from the Medicare Data on Provider Practice and Specialty (MDPPAS), a database of physicians with information on output, specialty, and practice location. I only include physicians classified as hospitalists³, internal medicine, general practice or family practice in at least one year. I drop physicians with consistently less than 70% of patients in a hospital setting, as the focus is on those who are affected by a hospital's decision to adopt. I use this data to construct the dependent variables in my analysis.

The CMS Shared Patient Data provides annual information on how often the same patient is seen by different providers, available from 2009 to 2015. For instance, when a primary care physician refers a patient to a specialist, both share that patient. I focus on shared patients between physician and hospital entities, specifically those seen on the same day, to identify physicians closely affiliated with hospitals. To refine the sample, I restrict to physician-hospital pairs with non-missing tax codes and at least 30 same-day shared patients per year. This helps exclude office-based primary care physicians and ensures consistency with physicians identified in the MDPPAS dataset. I then link hospitals within the pairs to two sources of EHR adoptions: the American Hospital Association (AHA) survey and the Healthcare Information and Management Systems Society (HIMSS) survey. I record the first year a hospital adopts an EHR, occurring if in AHA

³Physician who self-identifies as hospitalist or has 70% of patients in inpatient setting. I cannot distinguish between types of hospitalists as the data source manipulates this definition based on percent of patients seen in a hospital, so these can be formal hospitalists or internists.

the hospital “fully” uses an EHR, or in HIMSS the hospital lists the EHR vendor they employ. I focus on the initial shock of new technology and abstract away from some of the more granular characteristics of the technology for the purposes of this paper.

I use data from the AHA survey to measure hospital-physician integration, categorized into four levels from least to most integrated: independent practice association (IPA), open physician-hospital organization (OPHO), closed physician-hospital organization (CPHO), and integrated salary model (ISM). In an IPA, physicians maintain independence under loose contracts; OPHOs involve additional administrative duties; CPHOs include formal contracts with specific physicians; and ISMs employ physicians directly as salaried staff (Short, Ho, and McCracken, 2017). Finally, to maintain consistency, I exclude hospitalists affiliated with hospitals that were acquired during the study period.

I then aggregate the data to the hospitalist level. Since I only observe outcomes from 2016-2017 and not treatment, I drop anyone not treated by 2015. That is, the data contains no units which are “never-treated”. I limit to hospitalists who graduated before 2005, as anyone who graduated medical school after that will be finishing residency during the span of the data and will exhibit labor market changing behavior, which could be correlated to EHR exposure purely because of switching work settings during a time of rapid EHR implementation. There are various thresholds used in the construction of this final data, including the years chosen and how many shared patients constitutes a physician-hospital relationship. To ensure that the results are not sensitive to various choices of the thresholds, I estimate the effect of EHR exposure on all outcomes under different combinations of data limitations.

3.2 Summary Statistics

The first dependent variable I consider is the decision for a hospitalist to stop seeing Medicare patients, which I refer to as exit. Exit is an indicator variable equal to 1 in the first year that a hospitalist sees no patients in any future year. This captures a one-time decision to stop seeing patients, not a cumulative decision in all following years. The next variables I construct capture whether physicians change work setting. I consider the fraction of patients a physician sees in an office, and I construct an indicator variable equal to one if a physician sees any positive number of patients in an office. Finally, I consider outcome variables which explore care/billing patterns, the number of patients seen and claims filed per patient.

For each progression of decision making, I limit to hospitalists who did not change their behavior previously. That is, for office-based outcomes I exclude those who exit, and for care/billing outcomes I exclude those who exit or change their work setting. I show summary statistics of relevant variables based on each sample in Table 1.

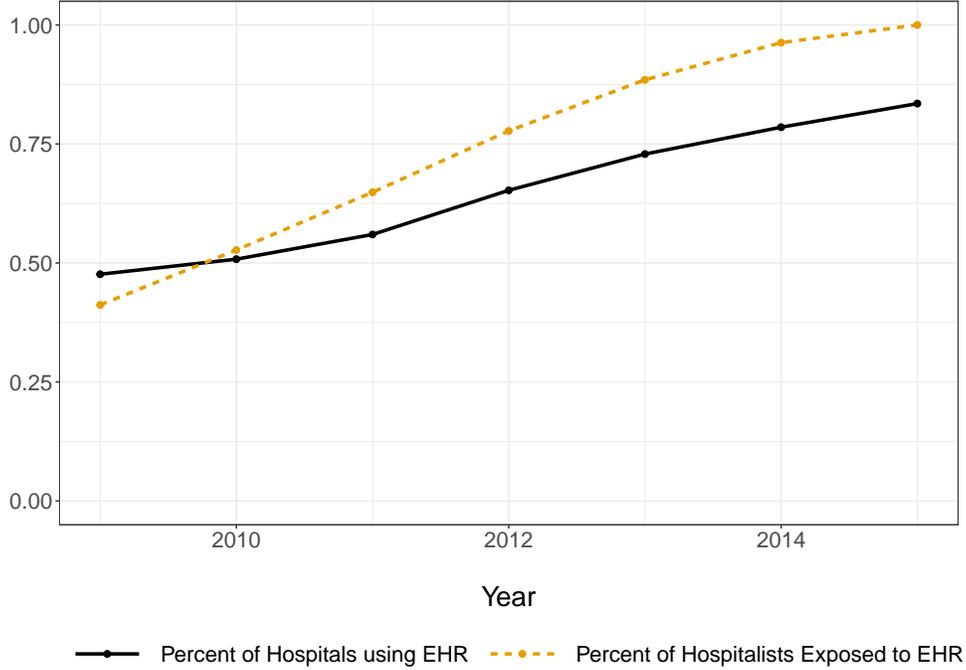
Table 1: Summary Statistics

Variable	Exit Sample		Office Sample		Productivity Sample	
	Mean	SD	Mean	SD	Mean	SD
Physician Characteristics						
Age	44.58	8.29	44.54	8.24	44.84	8.39
Female	0.33	0.47	0.33	0.47	0.34	0.47
Number of Hospitals Worked With	1.34	0.66	1.34	0.67	1.24	0.57
Number of Systems Worked With	1.19	0.46	1.19	0.46	1.12	0.37
Integration Characteristics						
Works w/ IPA Hosp.	0.26	0.44	0.26	0.44	0.24	0.43
Works w/ OPHO Hosp.	0.98	0.14	0.98	0.14	0.97	0.16
Works w/ CPHO Hosp.	0.14	0.35	0.14	0.35	0.12	0.32
Works w/ ISM Hosp.	0.68	0.47	0.68	0.46	0.65	0.48
EHR Exposure						
Year of EHR Exposure	2011	1.92	2011	1.92	2011	1.97
Outcome Variables						
Ever Exit Clinical Work	0.03	0.18				
Fraction Patients in Office			0.08	0.22		
Work in an Office			0.27	0.45		
Number of Patients					331	278
Claims per Patient					4.10	4.24
Num. Hospitalists	15,465		14,969		11,295	

I follow a maximum of 15,465 hospitalists over time, depending on the sample considered. Out of all hospitalists in the sample, 3% stop seeing patients at some point. Since this is a relatively rare event, I show a geographic distribution of exiters in Appendix A.1, which shows that exiting is not concentrated geographically. A positive number of hospitalists are exposed to an EHR in each year, with the majority of exposure is concentrated around 2011. For hospitalists in the office setting sample (those who do not exit), 27% of physicians have a positive number of patients seen in an office, and the average physician sees 8% of their total patient volume in an office. Hospitalists who do not exit or switch work setting see 331 patients and file 4 claims per patient on average. Consistent across all samples, hospitalists tend to work in 1.3 hospitals and 1.2 systems.

In Figure 2, I present a graph of the variation in treatment over time and the underlying hospital adoption that drives this treatment. In 2009, 41% of all hospitalists were already exposed to an EHR. That is, half of the sample had no affiliation with EHRs at the beginning of the sample period. Since I drop any physicians who do not have exposure to an EHR by 2015, 100% of the sample is exposed by the end of the sample period. In 2009, 47% of hospitals have adopted an

Figure 2: Treatment Variables Over Time



EHR, and this percentage grows steadily over time, driving the increase in exposed hospitalists.

4 Empirical Strategy

In this setting, treatment effects may be heterogeneous across hospitalists who are exposed in different years since the underlying characteristics of hospitals with different adoption timing may be correlated with the decision to adopt. Therefore, to avoid negative weighting issues this causes in a classical two way fixed effects specification, I instead estimate average treatment effects for a specific group g at time t :

$$ATT(g, t) = \mathbb{E}[Y_t(g) - Y_t(0) | G_g = 1],$$

where $G_g = 1$ for those in group g . A group indicates all hospitalists treated, or first exposed to an EHR, in the same year. I employ the doubly robust estimator established in Sant'Anna and Zhao (2020). Other estimators that similarly address the concerns yield similar results, presented in Appendix B.1. I aggregate the group time estimates to a more familiar event study plot with simultaneous confidence bands calculated using bootstrapped standard errors. Further, a weighted average of estimates is taken to provide a single ATT value for each outcome; these values are presented in the notes of each results plot.

For each outcome, I present aggregated group time treatment effects of being exposed to an EHR for the full sample of hospitalists, as well as split samples for those at least 60 (typical retirement age) and less than 60 years of age.⁴ The sub-samples are constructed based on maximum physician age throughout the sample, and thus remain fixed for each year.

4.1 Identification Assumptions

There are several assumptions necessary to identify the parameters of interest, $ATT(g, t)$. First, I assume that treatment is not reversed. That is, once a hospitalist is exposed to an EHR, they cannot be un-exposed. For hospitalists that do not switch hospitals, an EHR is a costly technology and requires a significant amount of collaboration to implement, so a hospital does not have incentive to un-implement an EHR. Any hospitalists who stop using EHRs through exiting or changing work settings are removed from the relevant sample, satisfying this assumption.

Second, I assume hospitalists do not anticipate EHR exposure prior to occurrence. If they learn that the system will be implemented and change behavior beforehand, the estimates may be biased. Generally, even a complex EHR system can be completely set up within a year, and most systems take 6 to 9 months (Uzialko, 2021). Therefore, I proceed in the main specification assuming no anticipation. However, I explore and present results for a one year anticipation period in the specification charts in Appendix B.4, and I find that accounting for one year of anticipation does not drastically change the findings.

As is usual in a difference-in-differences framework, I assume a version of parallel trends based on not-yet-treated units. I assume that average outcomes for those treated in group g would have followed a parallel trend to those in groups treated in later periods. While this assumption cannot be tested directly, I present p-values for a Wald test of pre-trends. Further, I investigate pre-trends extensively in Appendix B.5 by allowing for specified violations in the parallel trends assumption (Rambachan and Roth, 2019). My findings are robust to reasonable violations in the parallel trends assumption.

To interpret these estimators as the causal effect of EHR exposure on various labor market outcomes, there are additional institutional assumptions necessary. First, I assume that when a hospital adopts an EHR, hospitalists working in that hospital use the EHR. While some hospitalists may create work-arounds to ease the burden of EHR use, it is inevitable that they will need to interact with the technology in some capacity. Second, I assume that exposure to EHRs is exogenous. That is, hospitalists do not influence the hospital's decision to implement an EHR. Finally, I assume that

⁴Alternatively, I could estimate the effect for more age groups and observe the relationship between the overall ATT and age. However, the data does not support further sub-setting as sample sizes become small.

EHR adoption is not correlated with other hospital changes that could also affect physician work environment. To mitigate concerns along these lines, I present a hospital level analysis in Section 6.2, where I show that observable hospital behaviors are not associated with EHR adoption.

5 Effect of EHR Exposure on Labor Market Outcomes

5.1 No Longer Seeing Patients

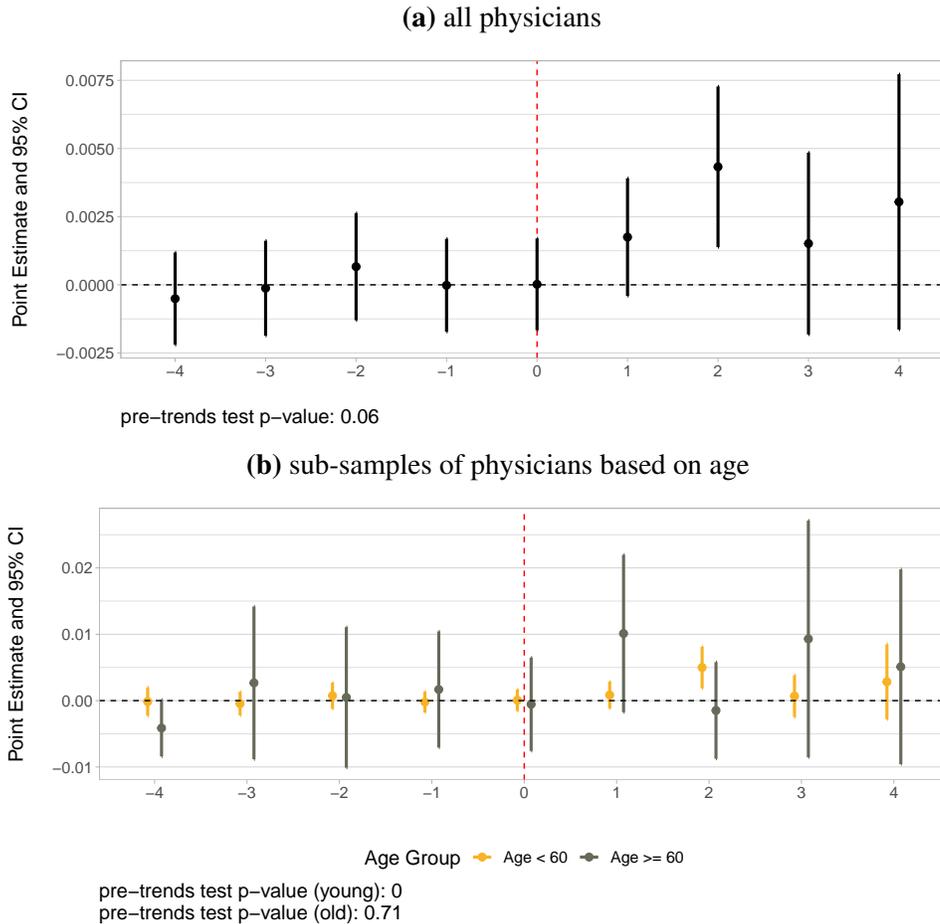
The effects that I present are a weighted average of the $ATT(g, t)$ parameters discussed in Section 4. Shown in Figure 3, being exposed to an EHR leads to a .2 ppt increase in the likelihood of exiting in the first year after exposure, and a .4 ppt increase in the second year after exposure. The coefficients remain similar in the years following, though noisier due to smaller sample of not-yet-treated groups. While numerically small, these effects are economically meaningful relative to the proportion of hospitalists who exit in the sample, .03. That is, the effect represents a 7% (13% in the second year) increase in the likelihood of exiting relative to the mean.

Next, I examine how the results differ when considering physicians in different age brackets: pre-retirement (≤ 60) and typical retirement age (> 60) in Figure 3b. Interestingly, older hospitalists are driving the positive estimate in the year after exposure, and younger hospitalists only exit in the second year after exposure. A retirement age hospitalist is 1 ppt more likely to exit after being exposed to an EHR, a 25% increase relative to the mean. In comparison, hospitalists less than 60 are not more likely to exit in the first year after exposure, but are approximately 16% more likely to exit two years after exposure. If younger hospitalists are leaving clinical work for administrative or consulting roles, the preparation time for this type of switch (preparing a resume, applying for new jobs, etc.) could be longer than one year, whereas formal retirement can be decided and executed within one year, which could explain this difference. Similarly, younger hospitalists may attempt to learn complex EHRs for a longer period of time before deciding the cost is too high.

5.2 Work Setting

I investigate another dimension by which hospitalists could change behavior after EHR exposure: the allocation of their time spent in different settings. Specifically, I measure whether the hospitalist, who begins with the majority of their patients seen in inpatient settings, shifts work to a stand-alone office setting. The outcomes I consider are (1) an indicator variable equal to 1 if the hospitalist has any positive number of patients in an office in a given year, and (2) the fraction of patients a hospitalist sees in an office in a given year. This analysis does not include individuals

Figure 3: Effect of EHR Exposure on No Longer Seeing Patients



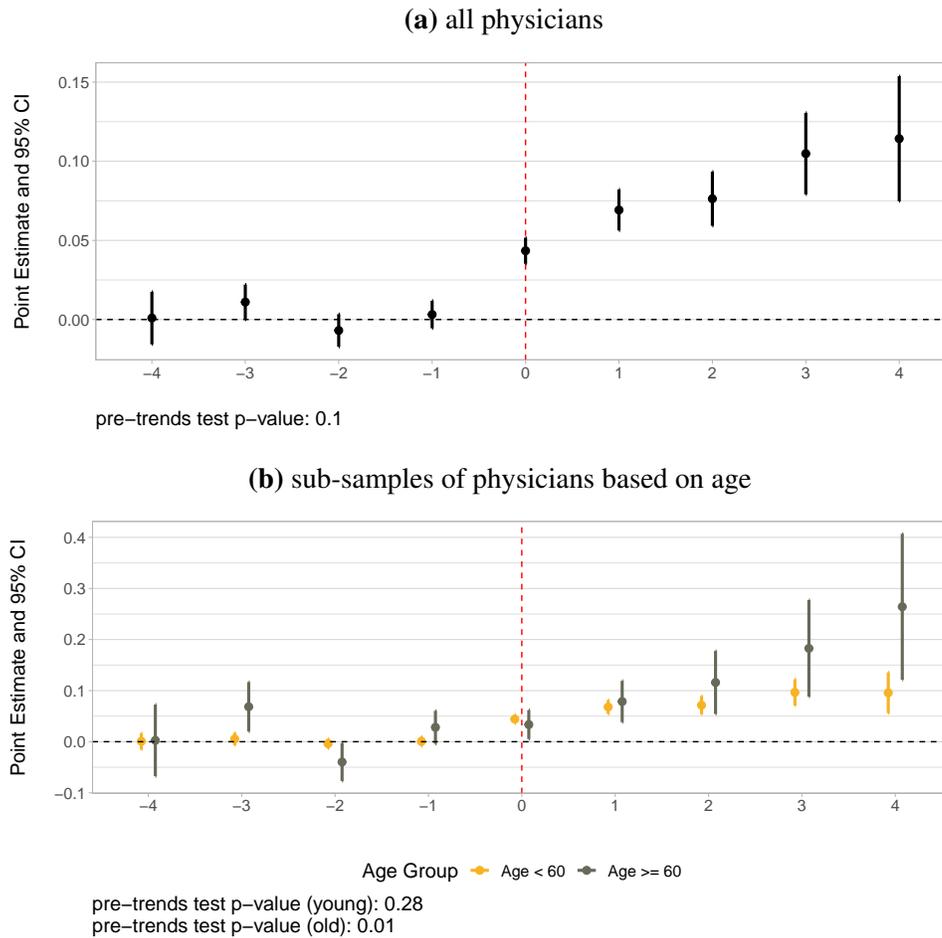
Notes: Panel (a) shows average group time treatment effects aggregated over groups to an event study plot. Panel (b) shows these results for different subgroups of physicians by age. The p-value listed for each graph corresponds to a Wald test for pre-trends. Confidence intervals shown are simultaneous confidence bands accounting for multiple hypothesis testing. ATT for all, < 60, >= 60 with SE in parentheses: 0.003 (0.0009), 0.002 (0.001), 0.006 (0.003), respectively.

who exit at any point in the sample, as a zero could indicate dropping out of the data instead of working solely in a hospital.

Shown in Figure 4a, EHR exposure increases the likelihood of working in an office by at least 4 ppts in every year after being exposed, where the effect increases over time (up to 10 ppts four years after exposure). This estimate is equivalent to a 15-38% increase relative to the average proportion of hospitalists who work in an office. That is, those working only in a hospital are more likely to start working in an office after being exposed to an EHR. However, unlike the estimated exiting effects, the increased likelihood of working in an office is persistent over time, which suggests that physicians do not move back to inpatient settings even after multiple years.

To continue investigating how different ages may respond differently to technology, I also

Figure 4: Effect of EHR Exposure on Likelihood of Working in Office



Notes: The top panel shows average group time treatment effects aggregated over groups to an event study plot. The bottom show these results for different subgroups of physicians by age. The p-value listed for each graph corresponds to a Wald test for pre-trends. Confidence intervals shown are simultaneous confidence bands accounting for multiple hypothesis testing. Overall ATT for all, < 60, >= 60 with SE in parentheses: .07 (0.008), .06 (0.008), 0.12 (0.03), respectively

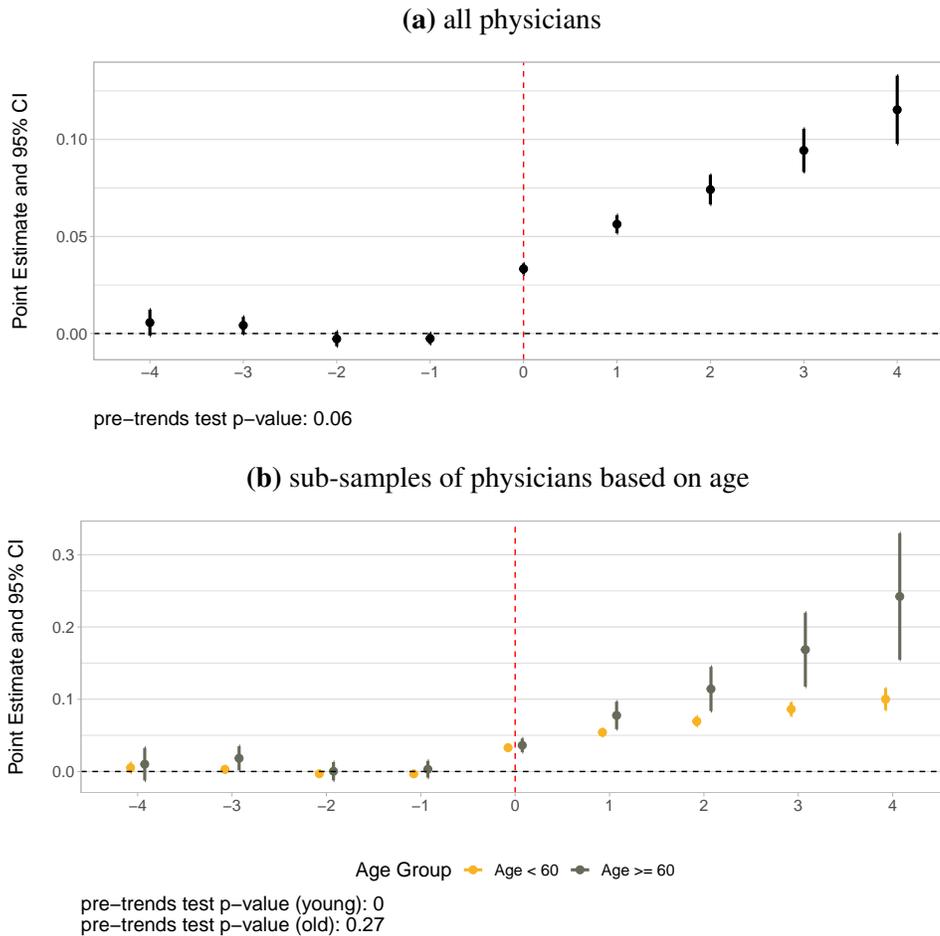
present estimates for hospitalists less than 60 and greater than or equal to 60 separately in Figure 4b. Unlike the choice to exit, both age groups respond to EHR exposure similarly. The magnitudes for senior hospitalists continue to increase over time, whereas the estimates for younger physicians become stagnant. Senior physicians are more likely to work in an office than younger physicians on average, so both groups exhibit similar increases relative to the mean.

I also estimate the effect of EHR exposure on the fraction of patients seen in office based settings.⁵ The estimates, displayed in Figure 5a, show an increase in the fraction of patients seen

⁵Since this variable is tied to the total number of patients seen, I also estimate this analysis with the outcome of

in office settings after the hospitalist is exposed to an EHR. In the year of exposure, the fraction of patients increases by .04, and by the fourth year after exposure the fraction of patients increases by .12. Since the hospitalists included in the sample are meant to have close ties with hospitals, the average fraction of patients seen in an office is low, at .09. Thus, the observed effect is large, equivalent to a 75-150% increase relative to the mean. Further, the variation in this variable can either be driven by physicians who already worked in an office or those who did not and then began working in an office after exposure.

Figure 5: Effect of EHR Exposure on Fraction of Patients Seen in Office



Notes: The top panel shows average group time treatment effects aggregated over groups to an event study plot. The bottom show these results for different subgroups of physicians by age. The p-value listed for each graph corresponds to a Wald test for pre-trends. Confidence intervals shown are simultaneous confidence bands accounting for multiple hypothesis testing. Overall ATT for all, < 60, >= 60 with SE in parentheses: .08 (0.003), .07 (0.003), 0.13 (0.01), respectively.

total number of patients seen in an office and the results are identical.

5.3 Patient Count and Billing Activity

I now estimate the effect of EHR exposure on patient count and billing activity per patient. I limit to hospitalists who work with the same hospital through the entire sample. These hospitalists do not exit or see patients in another setting, so that when an EHR is implemented they remain utilizing the EHR for the remainder of the sample. The effect of EHR exposure on patient count is shown in Figure 6a. For hospitalists of any age, patient count increases by anywhere from 20-55 patients after EHR exposure. These estimates translate to a 6-17% increase in the number of patients seen relative to the mean. Interestingly, when splitting the sample by age group, the effect decreases three and four years after exposure for younger hospitalists, while the effect for senior age hospitalists remains persistent over time, suggesting that technology's effect on patient care in the long run depends on age.

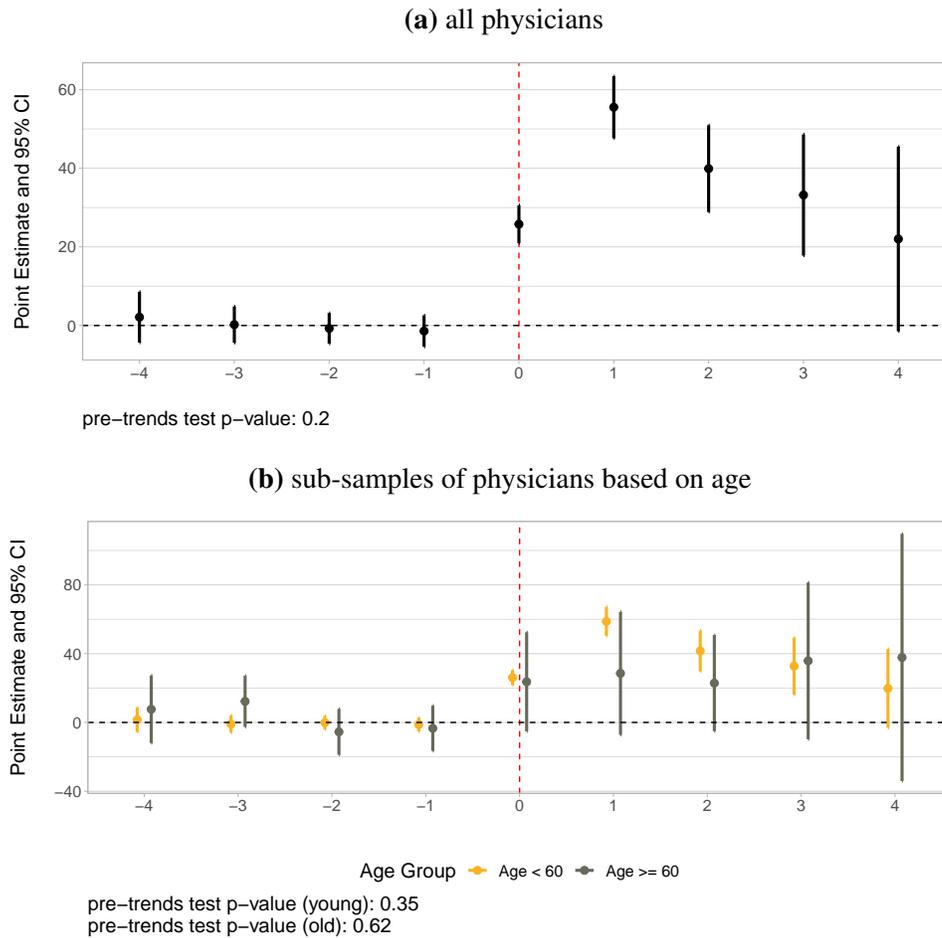
Finally, I consider whether EHR exposure affects claims filed per patient. On average, hospitalists file roughly 4 claims per patient. I present estimates in Figure 7a, which shows that EHRs tend to increase claims filed per patient over time. The magnitude of the effect ranges from .2 to .8, or 5-20% relative to the mean, indicating that as hospitalists are exposed to the new software, they bill for more items on average. I also show heterogeneity of this result by age group, and Figure 7b shows that younger hospitalists are exhibiting more of an increase in claim count than older hospitalists. However, when I analyze this relationship under different data specifications, shown in Appendix B.4, a majority of the estimates are negative or zero. Anticipation and limiting sample years cause the effect to disappear or even become negative. Thus, it is not clear whether this positive effect is actually an unintended consequence of EHRs, and I do not make any strong conclusions.

6 Alternative Explanations and Robustness Checks

6.1 Changes Related to the ACA

Many changes under the Affordable Care Act (ACA) from 2010–2014, such as insurance expansion and hospital quality initiatives, affected health care delivery. For these changes to bias my results, they would need to either drive hospital EHR adoption or differentially affect hospitalists exposed to EHRs versus those not exposed. While I cannot fully rule this out, it seems unlikely that any one policy caused large, immediate effects across all treated groups. If results were mainly driven by hospitals adopting EHRs in 2012 (pay-for-performance initiatives) or 2014 (Medicaid expansion), ACA policies could explain observed physician behavior changes. However, I present average group-time treatment effects by cohort in Appendix B.4, showing consistent patterns across groups regardless of EHR adoption year, suggesting ACA effects are not driving

Figure 6: Effect of EHR Exposure on Number of Patients Seen

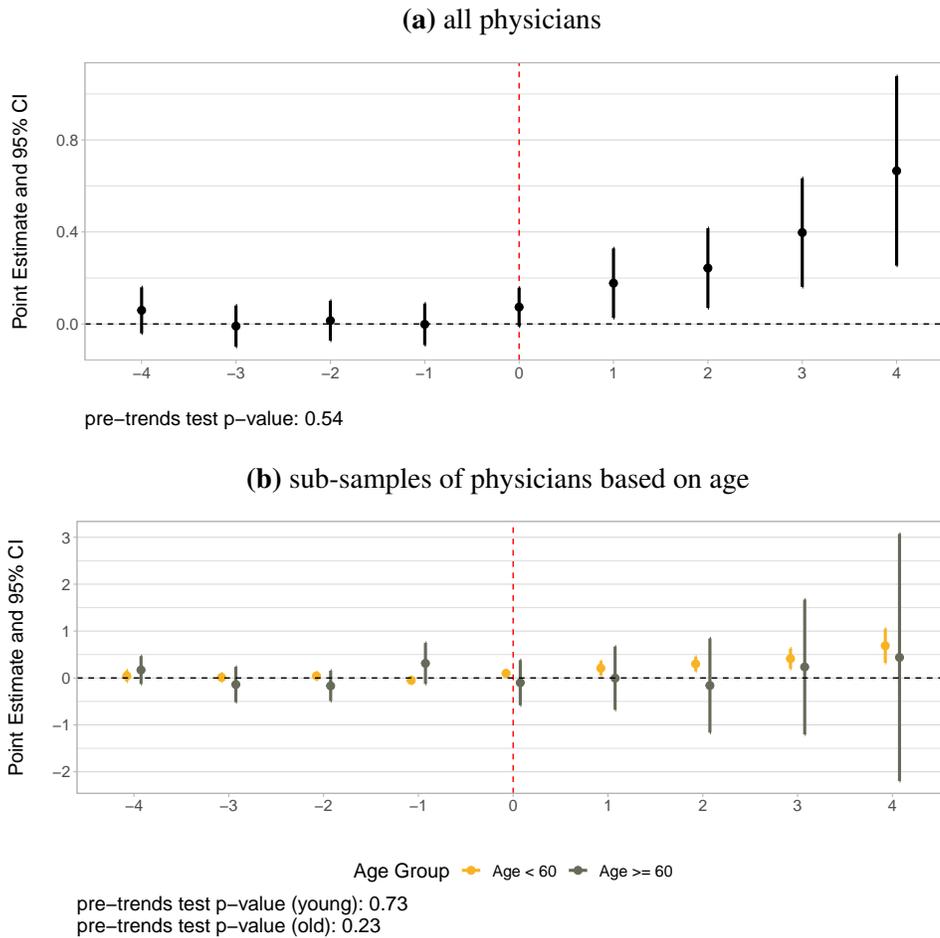


Notes: The top panel shows average group time treatment effects aggregated over groups to an event study plot. The bottom show these results for different subgroups of physicians by age. The p-value listed for each graph corresponds to a Wald test for pre-trends. Confidence intervals shown are simultaneous confidence bands accounting for multiple hypothesis testing. Overall ATT for all, < 60, >= 60 with SE in parentheses: 28 (5.0), 25.6 (5.2), 40.4 (15.0), respectively.

the main findings.

The ACA also had potentially opposing effects on Medicare patients: increased Medicaid demand could crowd them out, while small Medicare changes might boost their utilization. Yet recent evidence finds no negative spillovers onto Medicare patients (Carey, S. Miller, and Wherry, 2020). Another concern is that hospitalists exit Medicare practice to treat privately insured patients instead. However, EHR adoption is unlikely to correlate with Medicare coverage, hospitalists typically cannot select patients based on insurance, and no hospitalists in the sample opted out of Medicare according to CMS records.

Figure 7: Effect of EHR Exposure on Claims per Patient



Notes: The top panel shows average group time treatment effects aggregated over groups to an event study plot. The bottom show these results for different subgroups of physicians by age. The p-value listed for each graph corresponds to a Wald test for pre-trends. Confidence intervals shown are simultaneous confidence bands accounting for multiple hypothesis testing. Overall ATT for all, < 60, >= 60 with SE in parentheses: .33 (.09), .37 (.08), .05 (.44), respectively.

6.2 Other Hospital Changes

A potential concern is that EHR adoption coincides with other hospital initiatives to improve productivity or quality, which could bias the results if EHR adoption is not exogenous. I address this by examining whether hospital behaviors change alongside EHR adoption. If such changes exist, they likely offset some EHR costs, suggesting that my estimates for exit and work setting are lower bounds of the true effects.

Endogeneity is more relevant when considering patient outcomes. To investigate, I build a hospital-level dataset (2009–2015) and examine outcomes that might correlate with EHR adoption. I categorize these into two groups: (1) non-EHR investments and organizational changes (e.g.,

building costs, equipment costs, administrator turnover, bed counts), and (2) workforce changes (e.g., number of nurses and physicians, and levels of vertical integration with physicians). Expense data come from HCRIS; organizational characteristics from AHA.

Using the staggered difference-in-differences method from Callaway and Sant’Anna (2021), I estimate the association between EHR adoption and each outcome, with results shown in Table 2. I do not interpret these estimates causally but use them to assess concurrent changes. Across all measures, confidence intervals suggest no significant changes tied to EHR adoption. While I cannot rule out all unobserved shifts, the lack of observable changes mitigates concerns about confounding hospital initiatives.

Table 2: Association between EHR and Other Hosp. Behaviors

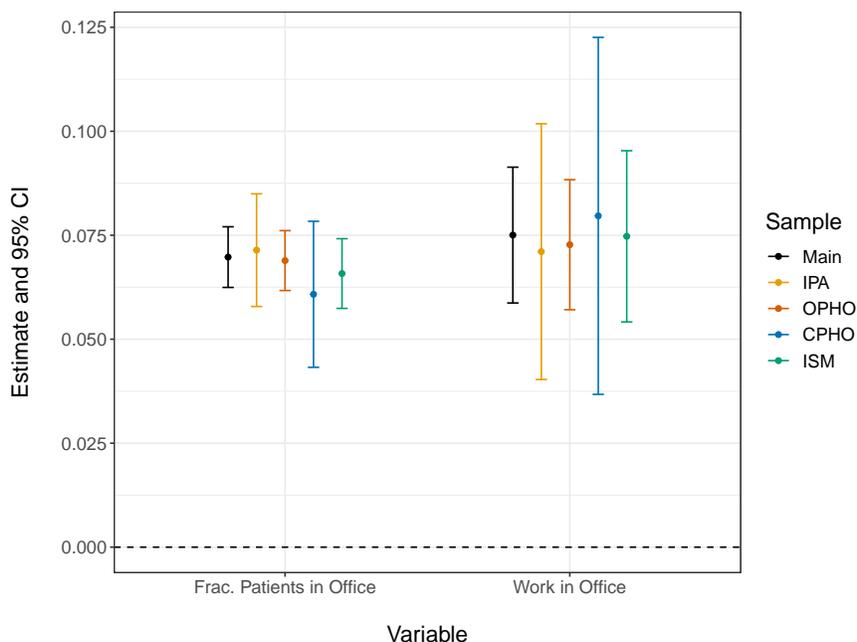
Outcome	Estimate	SE
Non-EHR Investment/Organizational Changes		
Capital Expenses (Buildings)	88,036	176,287
Capital Expenses (Equipment)	-96,304	184,225
Expenses w/out Med. Records (millions)	17.4	8.1
Owns Subsidiary	-0.02	0.02
Change Admin	0.01	0.02
No. Beds	0.28	2.15
Workforce		
No. Nurses	2.66	5.33
No. Physicians	0.00	0.01
IPA	-0.03	0.02
OPHO	0.03	0.01
CPHO	0.01	0.01
FIO	-0.04	0.02

6.3 Re-allocation of Hospitalists

As discussed in Section 2.2.2, hospitalists may shift work settings either by personal preference to avoid EHRs or due to hospital-driven reallocation enabled by EHRs. Although inpatient and outpatient EHRs differ, some substitution across settings is possible. To explore this mechanism, I conduct a heterogeneity analysis by hospital-physician contract type, which limits certain reallocation channels.

While I cannot observe individual physician contracts, I use hospital-level information on contract structures. I categorize hospitals by integration level: Independent Practice Association (IPA), Open Physician-Hospital Organization (OPHO), Closed Physician-Hospital Organization (CPHO), and Integrated Salary Model (ISM), ranging from least to most integrated. Hospitals are more

Figure 8: Effect of EHR Exposure on Outcomes by Integration Level



likely to shift physicians to outpatient settings when integration is high, whereas for less integrated hospitalists, reallocation is more voluntary.

The estimates, shown in Figure 8, do not vary meaningfully across integration types, suggesting that hospital-driven outpatient reallocation does not explain the main findings. Summary statistics and heterogeneity estimates are provided in Appendix B.2.

Another proxy for contract type is whether hospitalists work with one system or are connected to facilities from different systems, as a contracted hospitalist is more likely to work in multiple facilities and a salaried hospitalist is more likely to work in one facility. I present the results from this heterogeneity analysis in Appendix B.2, and find no difference in effects between the two groups. Finally, I show in Appendix B.3 that it is not likely that physicians are doing the same tasks but billing them in a different location.

6.4 Re-allocation of Patients

I investigate whether the increase in patients seen is driven by a decrease in the number of working hospitalists. Specifically, I test whether hospitals that lose hospitalists redistribute their patients among those remaining. While hospitals should eventually hire replacements, short-run estimates could reflect increased workloads rather than EHR effects, which may explain why patient counts spike initially and then level off over time. Using estimates from Section 5.1, I calculate the expected patient increase if all patients from exiting hospitalists were reassigned. In the year of

EHR exposure, there is no significant exit effect, so the increase in patient volume is attributable to EHR use. In the first year after exposure, patient increases from exits account for about 20 patients, less than the observed rise, suggesting EHR use is still the main driver initially. However, by the second year, exits become more prominent, and patient volume increases could partially reflect staffing reductions. Finally, Table 2 shows no correlation between EHR adoption and overall physician counts, indicating that hospitals replace exited physicians in the short run.

6.5 Evaluating Pre-Trends and Specification Charts

In Appendix B.5, I present confidence intervals robust to violations of the parallel trends assumption for each outcome. For exiting clinical settings, visual inspection shows no clear pre-trend violations, though a Wald test suggests some evidence among younger hospitalists. Because estimates for this group become null under minor trend violations, I focus on results for older hospitalists. For outcomes related to switching offices, both p-values and visual pre-trends suggest no violations, and robustness checks confirm the findings. Results for the number of patients seen are similarly robust, but estimates for claims per patient become null under any assumed violation, so I avoid drawing strong conclusions for that outcome.

I also test robustness across a range of specification changes, including allowing for anticipation, altering sample years, and adjusting thresholds (Section 3). Exit effects remain relatively stable and positive, though estimates grow noisier under stricter data limitations. Office-based work and patient volume results are consistently robust. However, estimates for claims per patient are often negative or zero across specifications, so I do not emphasize any effect on claims.

6.6 Lee Bounds for Sample Selection

EHR exposure affects which hospitalists are included in the analysis for both office work outcomes and patient count outcomes (since I drop those who exit or move), indicating a non-random attrition problem. That is, if productive or tech savvy hospitalists are the ones who do not change their behavior after EHR exposure, the estimates are biased. To investigate whether my results are robust to this issue, I construct average treatment effect bounds as in Lee (2009). For these bounds to be valid, a monotonicity assumption must hold on selection. In this case, EHR exposure must increase the likelihood of exiting or moving workplace settings and no hospitalists remain in the sample under treatment who would have exited or moved in the absence of EHR exposure. Since I find in earlier analyses that hospitalists are more likely to move and exit due to EHRs, this assumption is reasonable. I present the lower and upper bounds in Table 3, and they confirm that even if there is nonrandom sample selection, the qualitative results still hold.

Table 3: Lee Bounds for Selection

Variable	Lower	Upper
Work in Office	0.016	0.016
Fraction in Office	0.010	0.015
Num Patients	6.127	15.076
Claims per Patient	0.070	0.178

7 Conclusion

The adoption of health information technology—specifically complex electronic health record (EHR) systems—was expected to revolutionize the health care system by improving efficiency, reducing costs, and enhancing quality of care. However, these anticipated benefits have not been fully realized. A critical gap remains in understanding how physicians, the primary users of EHRs, respond to their rapid implementation. In this study, I examine the effect of hospital EHR adoption on various labor market behaviors of hospitalists, a group for whom EHR-related burdens have been a prominent concern linked to burnout. Physician labor markets have broad implications for access to care, quality of care, and health care costs.

Using data on shared patients, hospital EHR adoption, and physician billing activity, I employ a difference-in-differences framework to analyze this relationship. I find that hospitalists exposed to EHR adoption are more likely to exit clinical practice, with effects concentrated among physicians aged sixty and older. Additionally, EHR exposure increases the likelihood of switching work settings, with a notable shift from hospitals to outpatient environments. Among hospitalists who remain in their original setting, I find an increase in patient volume that cannot be fully explained by patient redistribution or selective physician attrition. These findings offer new insight into how technological change influences individual physician decision-making. Given growing physician shortages in the United States, future research should explore how technological innovation more broadly affects aggregate physician supply.

The results have several important implications. First, hospitalists’ willingness to incur switching costs to avoid EHR burdens suggests that physician behavior may directly counteract policy initiatives aimed at regulating practice or promoting technology adoption. Such changes may disrupt patient care through loss of continuity or reduced access. Policymakers must therefore anticipate behavioral responses when designing regulations that affect physician work environments. Furthermore, this study highlights the importance of considering the immediate short-run effects of technology adoption—a factor that both governments and hospitals should weigh carefully when implementing or incentivizing new health technologies.

References

- Agha, Leila (2014). “The effects of health information technology on the costs and quality of medical care”. In: *Journal of Health Economics* 34.1, pp. 19–30. ISSN: 01676296. DOI: 10.1016/j.jhealeco.2013.12.005. URL: <http://dx.doi.org/10.1016/j.jhealeco.2013.12.005>.
- Arndt, Brian G et al. (2017). “Tethered to the EHR: primary care physician workload assessment using EHR event log data and time-motion observations”. In: *The Annals of Family Medicine* 15.5, pp. 419–426.
- BLS (2015). *Hospitalist : Career outlook*. URL: <https://www.bls.gov/careeroutlook/2015/youre-a-what/hospitalist.htm#:~:text=The%5C%20Society%5C%20of%5C%20Hospital%5C%20Medicine%5C%20estimates%5C%20that%5C%20there%5C%20were%5C%20about,as%5C%20programs%5C%20continue%5C%20to%5C%20develop..>
- Bronsoler, Ari, Jr. Doyle Joseph J, and John Van Reenen (Sept. 2021). *The Impact of Healthcare IT on Clinical Quality, Productivity and Workers*. Working Paper 29218. National Bureau of Economic Research. DOI: 10.3386/w29218. URL: <http://www.nber.org/papers/w29218>.
- Buntin, Melinda Beuwkes et al. (2011). “The benefits of health information technology: A review of the recent literature shows predominantly positive results”. In: *Health Affairs* 30.3, pp. 464–471. ISSN: 02782715. DOI: 10.1377/hlthaff.2011.0178.
- Callaway, Brantly and Pedro HC Sant’Anna (2021). “Difference-in-differences with multiple time periods”. In: *Journal of Econometrics* 225.2, pp. 200–230.
- Carey, Colleen M, Sarah Miller, and Laura R Wherry (2020). “The impact of insurance expansions on the already insured: the Affordable Care Act and Medicare”. In: *American Economic Journal: Applied Economics* 12.4, pp. 288–318.
- Cengiz, Doruk et al. (2019). “The effect of minimum wages on low-wage jobs”. In: *The Quarterly Journal of Economics* 134.3, pp. 1405–1454.
- Chang, H., D. Childs, and A. Grayson (2009). *President Elect Urges Electronic Medical Records in 5 Years*. URL: <https://abcnews.go.com/Health/President44/story?id=6606536%5C&page=1>. (accessed: 03.09.2021).
- Civil Rights, Office for (2017). *HITECH Act Enforcement Interim Final Rule*. URL: <https://www.hhs.gov/hipaa/for-professionals/special-topics/hitech-act-enforcement-interim-final-rule/index.html>. (accessed: 03.09.2021).

- Collier, Roger (2017a). “Electronic Health Records Contributing to Physician Burnout”. In: *Canadian Medical Association Journal* 45.185, E1405–E1406. DOI: 10.1503/cmaj.109552.
- (2017b). *The challenges of physician retirement*.
- Czernik, Zuzanna et al. (2022). “Hospitalist Perceptions of Electronic Health Records: a Multi-site Survey”. In: *Journal of General Internal Medicine*, pp. 1–3.
- Dashboard, Health IT (2019). *Office-based Physician Electronic Health Record Adoption*. URL: <https://dashboard.healthit.gov/quickstats/pages/physician-ehr-adoption-trends.php>. (accessed: 03.10.2021).
- Dranove, David, Chris Forman, et al. (2014). “The trillion dollar conundrum: Complementarities and health information technology”. In: *American Economic Journal: Economic Policy* 6.4, pp. 239–70.
- Dranove, David, Craig Garthwaite, et al. (2015). “Investment subsidies and the adoption of electronic medical records in hospitals”. In: *Journal of health economics* 44, pp. 309–319.
- Dymek, Christine et al. (2021). “Building the evidence-base to reduce electronic health record-related clinician burden”. In: *Journal of the American Medical Informatics Association* 28.5, pp. 1057–1061.
- eCareers, Health (Nov. 2022). *The rise of hospitalists - and what comes next*. URL: <https://www.healthcareers.com/career-resources/healthcare-news/the-rise-of-hospitalists>.
- Evans, R Scott (2016). “Electronic health records: then, now, and in the future”. In: *Yearbook of medical informatics* 25.S 01, S48–S61.
- Freedman, Seth, Haizhen Lin, and Jeffrey Prince (2015). “Information Technology and Patient Health: Analyzing Outcomes, Populations, and Mechanisms”. In: *National Bureau of Economic Research*. ISSN: 0898-2937. DOI: 10.3386/w21389.
- Gardner, John (2021). *Two-stage differences in differences*. Tech. rep. Working paper.
- Gray, Bradley M et al. (2022). “Evolving practice choices by newly certified and more senior general internists: a cross-sectional and panel comparison”. In: *Annals of Internal Medicine* 175.7, pp. 1022–1027.
- Gustman, Alan L and Thomas L Steinmeier (1986). “A disaggregated, structural analysis of retirement by race, difficulty of work and health”. In: *The review of economics and statistics*, pp. 509–513.
- Healthco, Quatris. *Understanding Meaningful Use*. URL: <https://www.quatrishealthco.com/understanding-meaningful-use/>. (accessed: 03.12.2021).
- Hillestad, Richard et al. (2005). “Can electronic medical record systems transform health care? Potential health benefits, savings, and costs”. In: *Health affairs* 24.5, pp. 1103–1117.

- Hitt, Lorin M. and Prasanna Tambe (2016). “Health care information technology, work organization, and nursing home performance”. In: *Industrial and Labor Relations Review* 69.4, pp. 834–859. ISSN: 00197939. DOI: 10.1177/0019793916640493.
- Hospitalist, Today’s (June 2016). *How many hospitalists work outside the wards?* URL: <https://www.todayshospitalist.com/how-many-hospitalists-work-outside-the-wards/>.
- Lee, David S (2009). “Training, wages, and sample selection: Estimating sharp bounds on treatment effects”. In: *The Review of Economic Studies* 76.3, pp. 1071–1102.
- Loria, Keith (Nov. 2020). *Physicians leaving profession over Ehrs*. URL: <https://www.medicaleconomics.com/view/physicians-leaving-profession-over-ehrs>.
- Lyu, Heather et al. (2017). “Overtreatment in the united states”. In: *PloS one* 12.9, e0181970.
- McCullough, Jeffrey S., Stephen T. Parente, and Robert Town (2016). “Health information technology and patient outcomes: The role of information and labor coordination”. In: *RAND Journal of Economics* 47.1, pp. 207–236. ISSN: 17562171. DOI: 10.1111/1756-2171.12124.
- Messler, Jordan and Winthrop F Whitcomb (2015). “A history of the hospitalist movement”. In: *Obstetrics and Gynecology Clinics* 42.3, pp. 419–432.
- Meyerhoefer, Chad D. et al. (2016). “The Consequences of electronic health record adoption for physician productivity and birth outcomes”. In: *ILR Review* 69.4, pp. 860–889. ISSN: 00197939, 2162271X. URL: <https://www.jstor.org/stable/26726515>.
- Miller, Amalia R. and Catherine Tucker (2009). “Privacy protection and technology diffusion: The case of electronic medical records”. In: *Management Science* 55.7, pp. 1077–1093. ISSN: 00251909. DOI: 10.1287/mnsc.1090.1014.
- Rambachan, Ashesh and Jonathan Roth (2019). “An honest approach to parallel trends”. In: *Unpublished manuscript, Harvard University*.
- Ringel, Marc (May 2019). *Electronic Health Records and doctor burnout*. URL: <https://blogs.scientificamerican.com/observations/electronic-health-records-and-doctor-burnout/>.
- Sant’Anna, Pedro HC and Jun Zhao (2020). “Doubly robust difference-in-differences estimators”. In: *Journal of Econometrics* 219.1, pp. 101–122.
- Short, Marah, Vivian Ho, and Ayse McCracken (2017). “The integration and de-integration of physicians and hospitals over time”. In: *Baker Institute for Public Policy, Houston, TX*.
- Stock, James H, David A Wise, et al. (1990). “The pension inducement to retire: An option value analysis”. In: *Issues in the Economics of Aging*, pp. 205–224.
- Tipping, Matthew D et al. (2010). “Where did the day go?—A time-motion study of hospitalists”. In: *Journal of hospital medicine* 5.6, pp. 323–328.

Uzialko, Adam (Feb. 2021). *How to implement an electronic health records system*. URL: <https://www.businessnewsdaily.com/9102-implement-an-electronic-health-records-system.html>.

A Data

I begin by connecting every National Provider Identifier (NPI)⁶, an identifier of each medical entity, to a description of their tax code⁷. I then categorize NPIs by key words in their tax code description. Any description containing “Internal Medicine”, “Hospitalist”, “Family Medicine”, or “General Practice” is classified as a primary care physician (PCP), and any description containing “hospital” is classified as a hospital. I create two data sets, one containing only PCPs and one containing only hospitals. These data are then used in identifying the relevant NPIs in the CMS Shared Patient Data.

For years 2009-2015, the CMS collected detailed information on the number of Medicare patients shared between any two NPIs within 30, 90, or 180 day intervals.⁸ In this paper, I use the data files that capture shared patient activity in 30 day intervals. I merge the shared patient data to the filtered tax code data to identify NPIs who are either PCPs or hospitals. I limit to pairs including one physician and one hospital. I combine duplicates into one observation, summing the same day count variable. Most of the pairs have very few shared patients, which is not indicative of the physician working inside the hospital. I drop any pairs who do not have at least 30 same-day shared patients per year the pair appears in the data.⁹

I combine the physician-hospital pairs with various data sets containing hospital and/or physician level information. First, I merge to CMS Physician Compare¹⁰ for information on each physician’s graduation date. This data contains more information on physician quality, but I am limited to time-invariant information since Physician Compare spans 2012-2015 and 2010-2012 is a time of major EHR implementation. I drop any physicians who graduated medical school after 2004, since graduating after that means leaving residency or graduating during the span of the main data, 2009-2017, and potentially exhibiting labor market changes that seem associated with EHRs but are not.

For hospital level variables, I use an AHA-NPI crosswalk to merge the pairs to the Annual Hospital Administration Survey (AHA Survey) from 2009-2015. Further, I use a hospital’s Medicare ID number to merge this data to the Healthcare Information and Management Systems Society (HIMSS) data. From these data sets I collect information on each hospital’s EHR use. I define EHR use as either answering in the AHA Survey that the hospital uses an EHR fully, or reporting which vendor is used in the HIMSS Survey. Next, I drop hospitals with missing information for

⁶<https://download.cms.gov/nppes/NPI-Files.html>

⁷<https://nucc.org/index.php/code-sets-mainmenu-41/provider-taxonomy-mainmenu-40/pdf-mainmenu-53>

⁸<https://www.nber.org/research/data/physician-shared-patient-patterns-data>

⁹I also consider thresholds of 10 and 60. Results are similar and can be found in Figure B.4.

¹⁰<https://data.cms.gov/provider-data/dataset/mj5m-pzi6>

EHR use. If a hospital does not answer the survey question in one year, but the year before and after have an identical survey answer, I fill in the missing year of information. I make a further limitation to hospitals with at least 10 beds. I sum a physician's same day count over their hospitals in a given year to create a physician level patient count variable.

The final data contains year, physician NPI, graduation year, years of experience, number of hospitals, minimum year exposed to EHR, number of systems, and an indicator for whether the physician changes hospitals. I complete this data to include years 2016 and 2017, but leave time varying variables missing for those years. Thus, this is a balanced panel of physicians.

A.1 Dependent Variables

Using physician NPI, I merge the physician treatment data to Medicare Data on Provider Practice and Specialty (MD-PPAS)¹¹, which spans 2009-2017. This data contains variables on physician specialty, Medicare claim counts in various zip codes, unique number of patients seen, fraction of patients seen in specific settings, patient demographics, and physician date of birth. Once the data is merged, I make a further limitation to drop any physicians with less than 70% of their total patients in a hospital setting.¹²

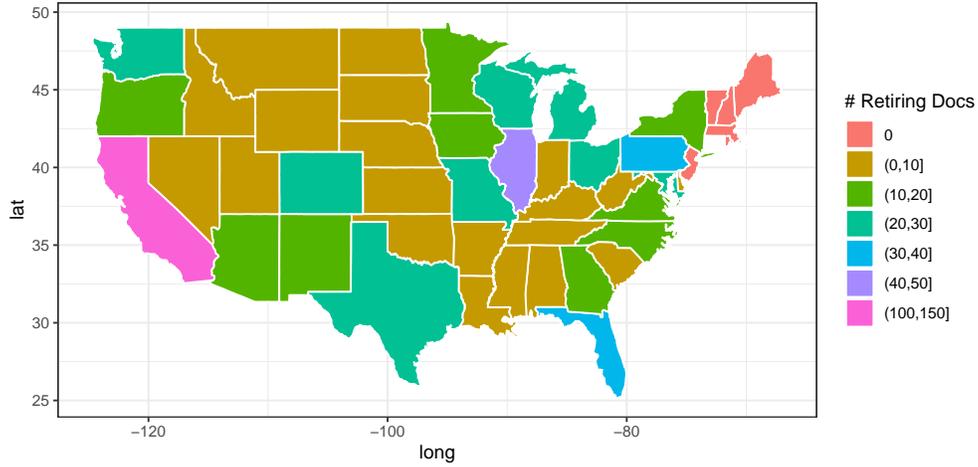
Next, I create the dependent variables used in the analysis. First, I sum a physician's claim counts across zip codes into one variable for total claim count in a given year. Then, I create a variable that sums up a physicians claims in all future years. In 2009, this variable sums up all claims from 2010-2017, and so on. Then I create a variable for the first year that a physician has a future claim count of zero. However, this counts exit one year too early. For example, if a physician has no claims from 2014 to 2017, the minimum year variable is set to 2013, but the first year of exit is 2014. Thus, I add one year to the minimum year variable. I also create a physician level variable for whether they ever retire for sample limitations in analyzing other outcomes. Since exit is relatively rare, I present a distribution showing where they occur geographically in Figure 9. The distribution is fairly uniform.

The next outcome variables consider level of labor market activity in office settings. I set all missing claim counts to zero. The data already contains a variable for the fraction of patients seen in an office setting, which is used in the analysis, but I further create an indicator variable for whether the hospitalist sees any positive amount of patients in an office setting. Finally, the variables for unique patient count and claim count are already defined.

¹¹<https://resdac.org/cms-data/files/md-ppas>

¹²I consider different thresholds as a sensitivity analysis. The results are similar, and can be found in Section B.4

Figure 9: Distribution of Exitters



B Sensitivity Analysis

B.1 Alternative Estimators

There is a robust recent literature pointing out issues with estimating staggered difference-in-differences with two-way fixed effects. To avoid these problems, I use the average group time effect estimator established in Callaway and Sant’Anna (2021). In this section, I discuss other potential estimators and why I ultimately included average group time treatment effects as the main specification.

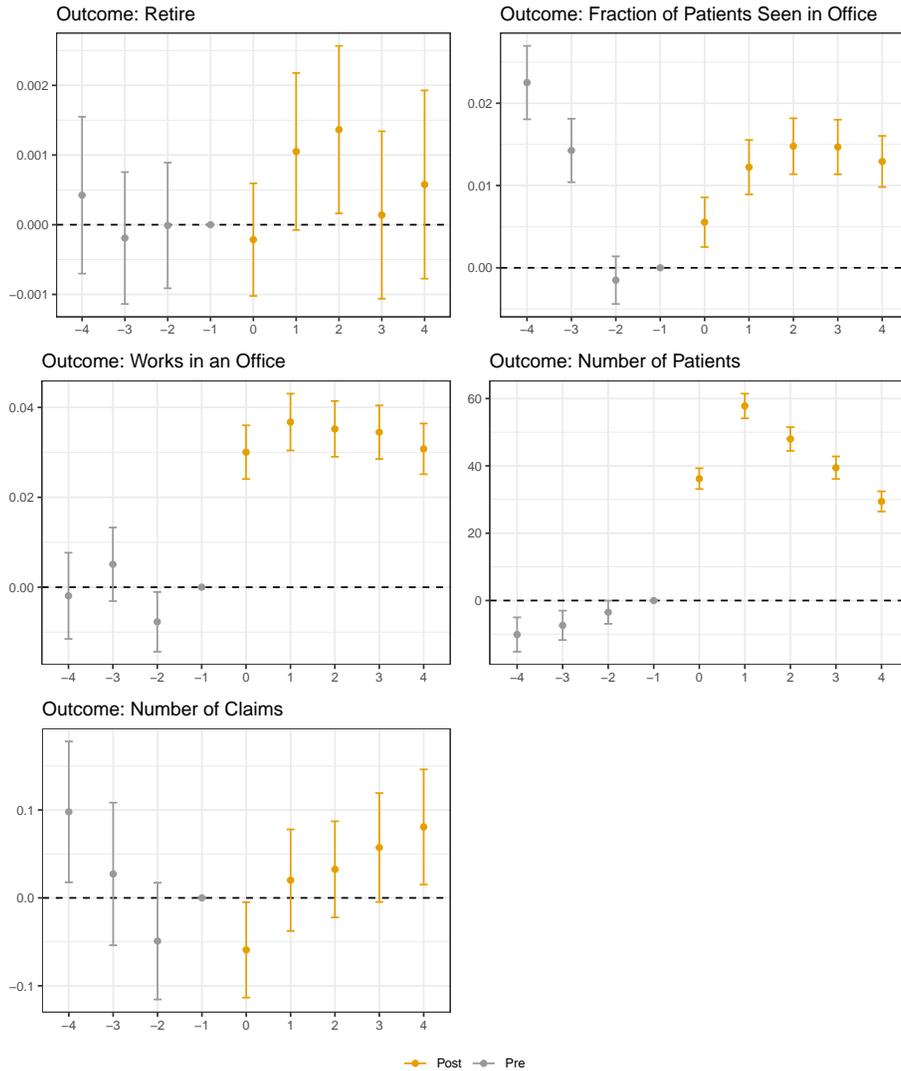
Before I present these estimators, I present results from a two way fixed effects setup that may suffer bias from negative weighting. Specifically, I estimate the following equation:

$$\text{outcome}_{i,t} = \sum_{j=-4}^{-2} \text{rel_year}_j + \sum_{j=0}^4 \text{rel_year}_j + \delta_i + \gamma_t,$$

where $\text{outcome}_{i,t}$ is one of the five outcomes discussed previously, rel_year_j is an indicator equal to 1 if year t is j years relative to treatment, δ_i are physician fixed effects, and γ_t are year fixed effects. These results from each outcome are presented in Figure 10. These results show similar patterns to the main specification, with slight variations in magnitude and more indications of a pre-trend.

One method that addresses issues with the TWFE specification by differencing out fixed effects is established in Gardner (2021). This approach is commonly known as two stage difference-in-

Figure 10: Results: Two Way Fixed Effects

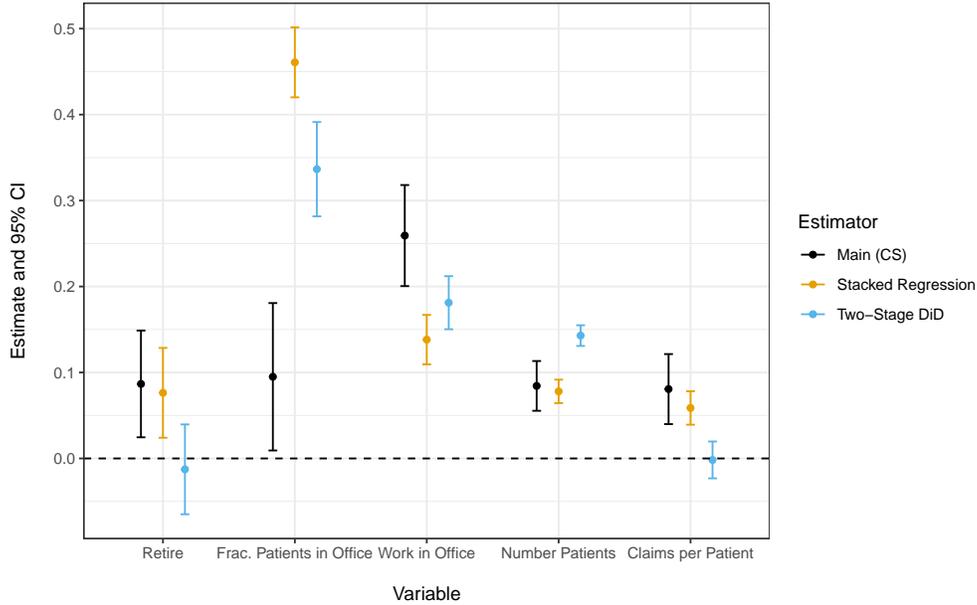


Notes: Event study plots from two way fixed effects estimation are shown for each outcome variable.

differences. The approach requires a group of never-treated units as the comparison group; a drawback of this study in that I remove all never-treated units to utilize all years of MD-PPAS data, and because never-treated hospitals are likely compositionally different from those treated in the majority of the sample. Nevertheless, I limit the years of data to 2009-2015 and present estimates as a percentage of the mean of each variable in Figure 11. Generally, the results have the same sign with a slightly larger magnitude.

Another estimation method first used by Cengiz et al. (2019) is commonly known as stacked regression. In this method, one re-frames event study data into groups of sub-experiments that

Figure 11: Alternative Estimators



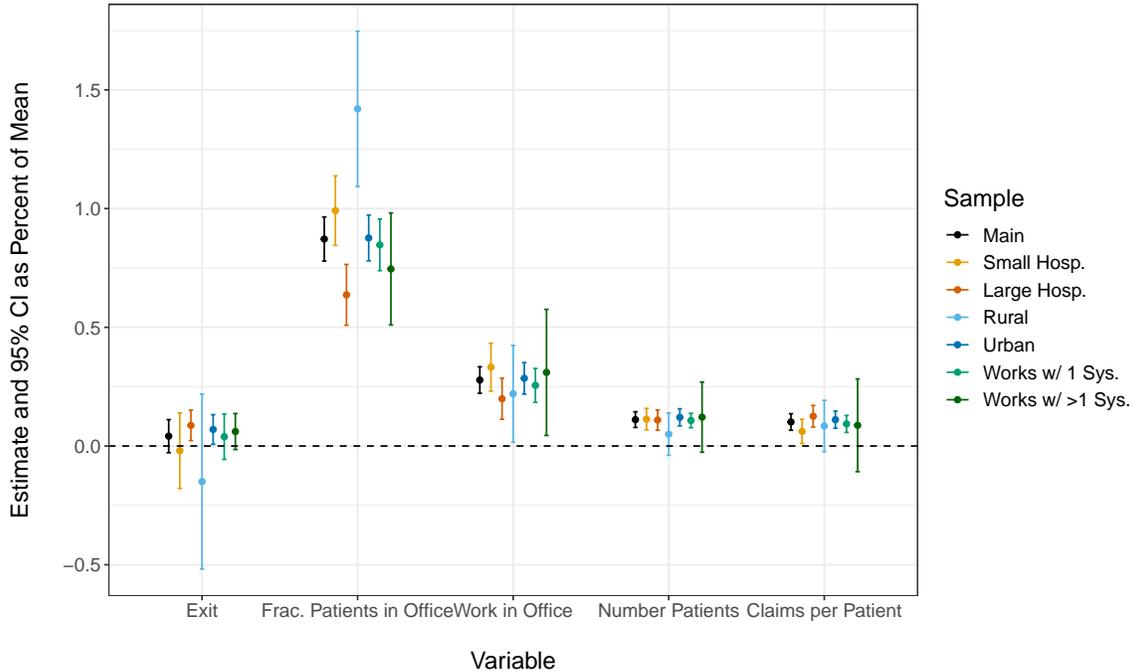
are then stacked on top of each other, defining a group of treated units and “clean” control units. The drawback of this method is choosing an event study window that must be the same for each treatment group, where a large window eliminates many treated units and a small window does not allow for estimation over many years. I present the results using this method with an event window of one year in Figure 11. The results are very similar to the main specification, with a smaller magnitude in some cases.

B.2 Heterogeneity

The main analysis focuses on the differential effects of technology for older vs. younger physicians. However, another important question is whether physicians in rural vs. urban settings respond differently to the technology since the implications for behavior changes are high in areas that already suffer from access to care issues. Thus, I present average treatment effects for each outcome where the sample is split between physicians in rural and urban areas. Further, to investigate the same idea, I present results for physicians who typically work at small hospitals vs. large hospitals. These average treatment effects as a percentage of the mean of each variable are shown in Figure 12. The results do not differ along these dimensions, implying that a new technology affects hospitalists across many different settings.

I also present a heterogeneity analysis for all outcomes broken down by level of integration the physician is exposed to in their connected hospitals. From least integrated (contracts with multiple facilities) to most integrated (full salaried hospitalists), results are broken down by Independent

Figure 12: Heterogeneity Analysis



Practice Association (IPA), Open Physician-Hospital Organization (OPHO), Closed Physician-Hospital Organization (CPHO), and Integrated Salary Model (ISM). Estimates do not differ between the types of integration relationships. I present summary statistics by contract type and estimates of this heterogeneity analysis for all outcomes.

B.3 Coding Activity as Cause for Work Setting Shift

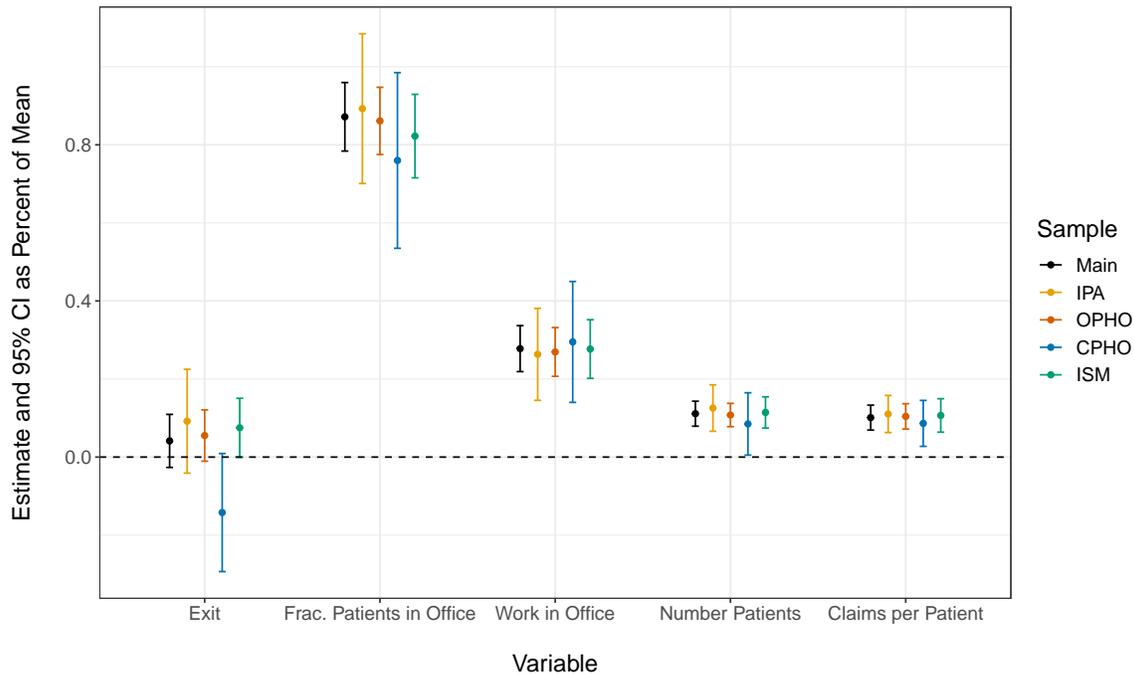
In this section, I investigate specific codes that take place in offices vs. inpatient settings to see what might be driving the increased office work. I merge the CMS Public Utilization File on Providers by Service (PUF), which provides data beginning in 2013. Therefore, for the purpose of this exercise, I consider only years 2013-2017. Using Provider NPI, I merge in the total number of claims and the number of claims that are specific to primary care services in hospitals and offices. These are listed on the y-axis of Figure 14. Then, using the same difference-in-differences methodology as the main results, I estimate the effect of EHR exposure on physician fraction of claims in each category.

The estimates and 95% CI are shown in Figure 14. After EHR exposure, the entire substitution from hospital to outpatient services is driven by a decrease in subsequent care in the hospital and an increase in established patient visits in an office. Subsequent visit in a hospital refers to evaluation/management of a patient during their hospital admission, while an established patient visit in

Table 4: Summary Statistics by Integration Type

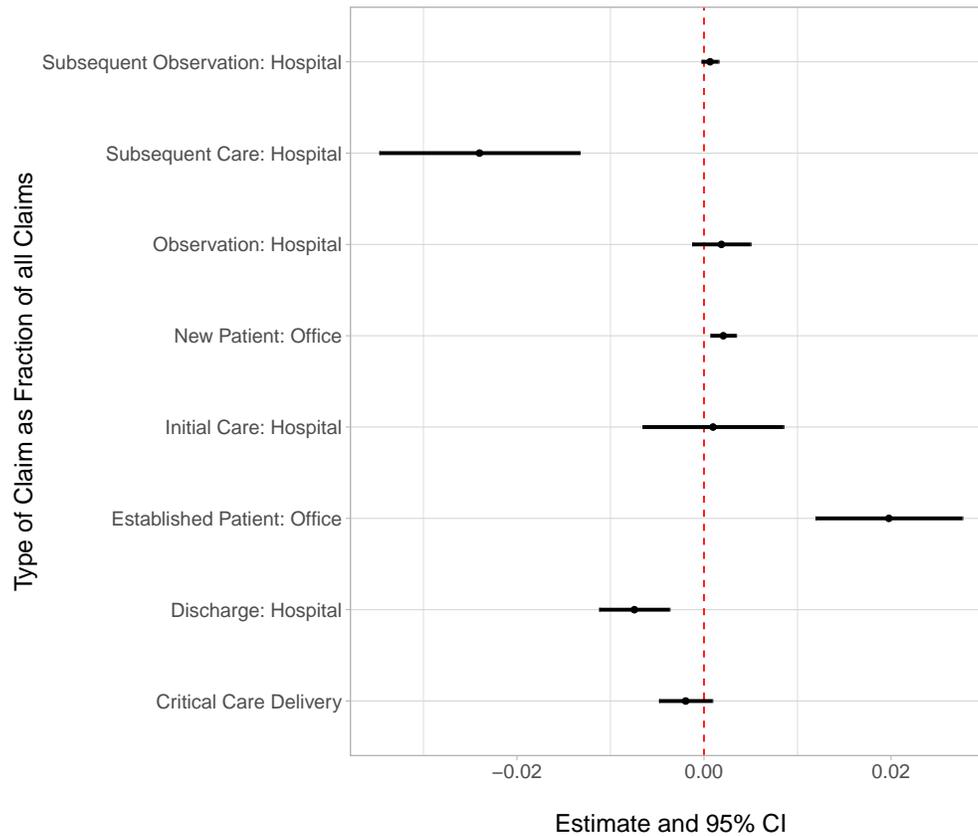
Variable	IPA		OPHO		CPHO		FIO		
	Mean	s.d.	Mean	s.d.	Mean	s.d.	Mean	s.d.	
Characteristics									
Age	44.69	8.30	44.56	8.28	44.27	8.18	44.43	8.18	
Female	0.34	0.47	0.33	0.47	0.34	0.47	0.33	0.47	
No. of Hospitals	1.42	0.74	1.35	0.67	1.48	0.78	1.38	0.70	
No. of Systems	1.25	0.53	1.19	0.46	1.29	0.56	1.22	0.48	
Year of EHR Exposure	2010.8	1.92	2010.9	1.91	2010.8	1.88	2010.8	1.88	
Outcomes									
Ever Exit Clinical Work	0.03	0.16	0.03	0.18	0.03	0.17	0.03	0.17	
Frac. Patients in Office	0.08	0.21	0.08	0.21	0.08	0.21	0.08	0.22	
Work in an Office	0.27	0.44	0.27	0.44	0.26	0.44	0.27	0.44	
Number of Patients	331	282	332	276	351	286	333	272	
Claims per Patient	3.99	3.47	4.04	4.09	3.87	3.44	3.97	4.34	
No. Hospitalists	3,994		15,152		2,172		10,550		

Figure 13: Heterogeneity Analysis, Hospital Contract Types



an office is when a patient who has been seen by the physician in the last three years returns with a new or existing condition. While both services can be performed by physicians who specialize in primary care, they are not substitutes. Specifically, codes for subsequent hospital inpatient or observation care codes (99231-99233) represent services that occur after the first encounter of the

Figure 14: Effect of EHR Exposure on Type of Claims



patient’s hospital admission, while codes for established patient office visit (99211 - 99215) must be in an office setting.

B.4 Various Changes to Specification

For each outcome, there are various adjustments I make to the main specification in order to show robustness. Some of these changes address potential endogeneity concerns, while some provide evidence that the results do not depend on thresholds which may be arbitrary. I will outline each of these adjustments and then present specification charts which account for different combinations of these changes.

While there is evidence to suggest the time it takes to implement an EHR does not exceed one year, there still may be some anticipation if physicians are expecting implementation in the hospitals they work in. Thus, the specification charts that I present in Section B.4 include specifications where one year of anticipation occurs.

In the main sample, I define dependent variables using years 2009 to 2017 and drop any hospitalists who were not exposed to an EHR by 2015 due to data constraints. This leads to estimation

based on not yet treated units, because all never treated units are dropped. As a robustness check, I limit sample years to 2009-2015 and use never treated units as the control group.

There are two main thresholds I define in the data process that were arbitrarily chosen. Thus, I present alternate specifications that adjust these thresholds. First, when deciding which hospitals a physician is working closely with, I drop any pairs that fall below 30 shared patients per year. The goal of this threshold is to eliminate pairs for which the hospitalist does not have close ties to the hospital. In alternate specifications, I lower the threshold to 10 patients per year and raise the threshold to 60 patients per year. Second, I only want to include hospitalists in the analysis, and thus drop any physicians that fall below 70% of patients in an inpatient setting. In alternate specifications, I change this threshold to 50% and 90%. While these are still arbitrary, they provide some evidence that changing the threshold does not dramatically change the results. I present overall average treatment effects for each outcome variable under different combinations of the specification adjustments. On each chart, the main specification is labeled in green and the average confidence bands across all combinations is shown in orange.

B.4.1 Leaving Clinical Setting

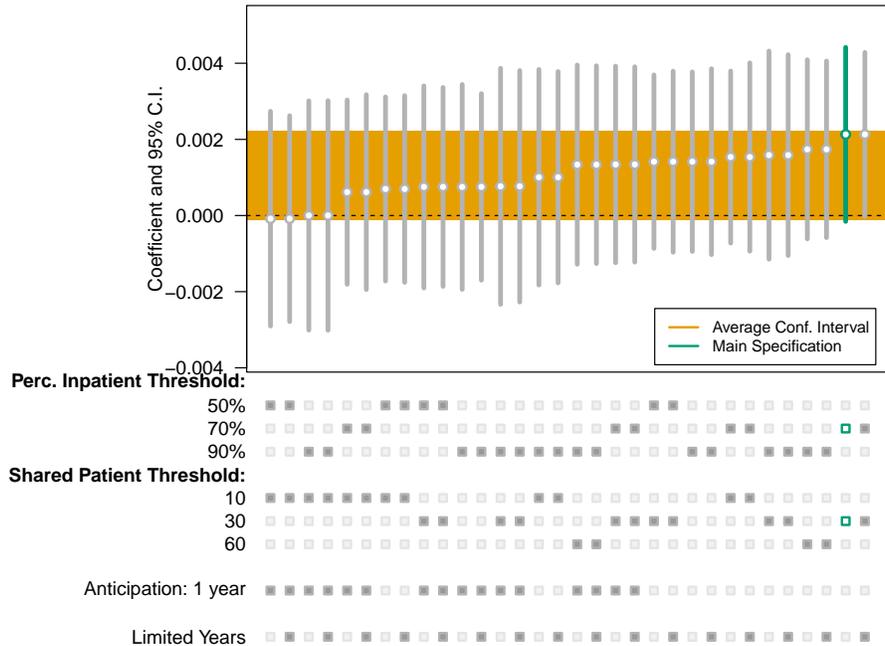
The effect of EHR exposure on leaving clinical work under different data specifications is presented in Figure 15. The main specification yields a marginally significant, positive average treatment effect of .0025. This coefficient is on the higher end of coefficients from all combinations of specifications. However, every data choice leads to a non-negative treatment effect. While it is possible that the true effect is zero, a likely reason for such noisy estimates is the rarity of exiting in the sample combined with more restrictive data choices. The average confidence band supports a positive effect of EHR exposure on leaving clinical settings.

B.4.2 Change Work Setting

The specification charts for outcomes regarding switching to outpatient settings are presented in Figures 17 and 18, respectively. First, I discuss the overall average treatment effect of EHR exposure on the likelihood of working in an office. The main specification yields finding that EHR exposure leads to a .08 ppt increase in the likelihood of working in an office. This is on the high end of all combinations of changes to specification, but all reveal a positive and significant effect ranging from .03 to .09.

Now I discuss the results for whether hospitalists change the fraction of patients seen in an office setting due to EHR exposure. Similarly, the ATT from the main specification is on the high end of other ATTs found, but all estimates suggest a positive effect. The average of all the confidence intervals found is (.04, .055), suggesting a positive effect on fraction of patients seen in

Figure 15: Outcome: Exit Clinical Setting



an office. Changes that affect the magnitude of the estimate (but not the sign) are limiting sample years and anticipation. Limiting the years in the sample and allowing for one year of anticipation both push the magnitude of the estimate down, but do not change the sign or significance.

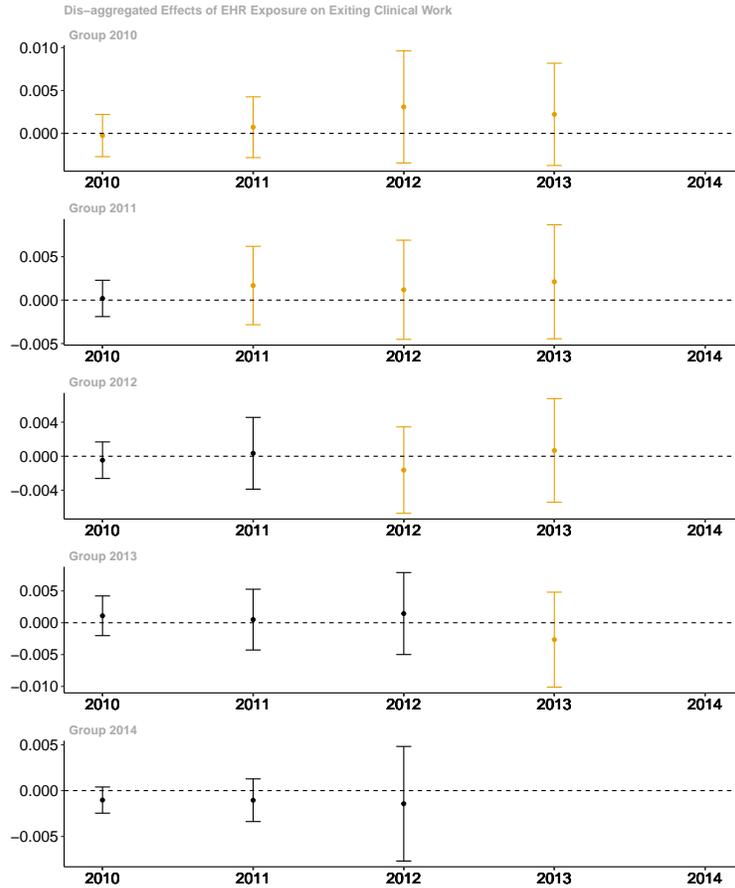
B.4.3 Patient Count and Claims per Patient

In Figure 21, I present estimates and confidence intervals for different combinations of specification changes. The result from the main specification is that EHR exposure increases patient count by 35. Under any data or specification adjustments, this estimate is nearly identical.

Lastly, I examine a chart of estimates for different specifications attempting to capture the effect of EHR exposure on claims per patient. My main finding is that EHR exposure increases claims per patient by around .3. However, various robustness checks reveal that this is overestimated, and a third of the various specifications yield negative coefficients. Specifically, anticipation and limiting sample years seem to flip the sign. Since the main finding is not consistent across various specifications, I make no conclusion that EHR exposure affects claims per patient.

As referenced in the main paper, the effect of EHR exposure on patient count could be biased under any demand-driven events happening at the same time as EHR adoption in hospitals. The only nation-wide demand altering event that I know of during the time of this sample is the Affordable Care Act, which took effect mainly in 2014. If this is driving the positive finding, this would show up in a de-aggregated treatment effects by year of exposure. In Figure 23, I show that the

Figure 16: Effect of EHR Exposure on Exiting Clinical Setting by Group



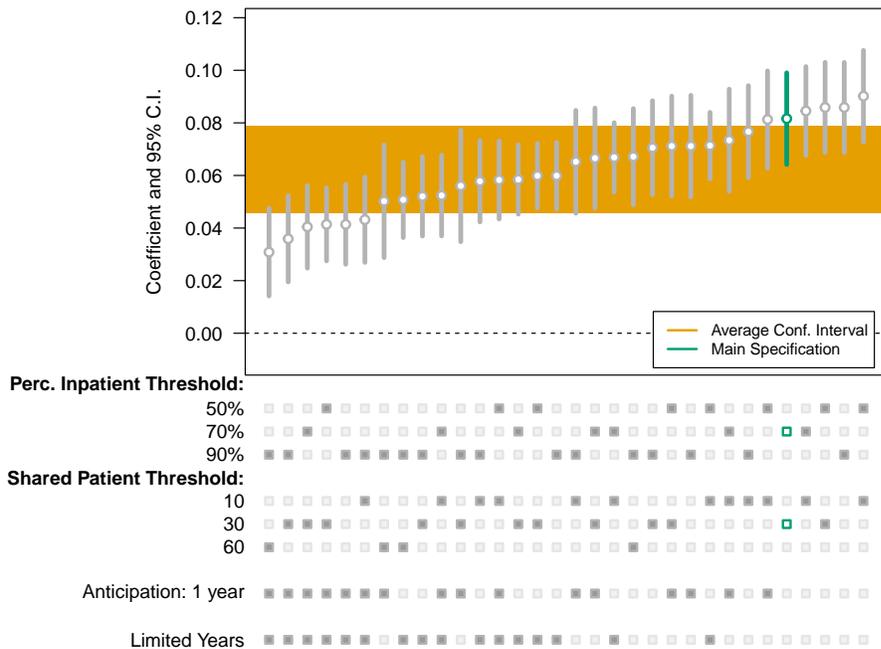
Notes: Plot shows average group time treatment effects for each treatment group.

effects persist in all years of exposure, not just in 2014.

B.5 Parallel Trends Assumption

Work by Rambachan and Roth (2019) points to the need to carefully assess the parallel trends assumptions commonly used in difference-in-differences. This is relevant in my analysis in the context of potential endogeneity. In my main specification I assume that, conditional on covariates, average outcomes for those treated in group g would have followed a parallel trend as those in groups treated in later periods in the absence of treatment. I assess this assumption by testing whether the coefficients showing the effect of EHR exposure on an outcome prior to actual treatment are jointly zero. All main specification graphs show p-values from this test. In most analyses, I fail to reject the null that there does not exist a pre-trend. Yet, in some cases, I reject this null. This section utilizes the tools established by Rambachan and Roth (2019) to demonstrate robustness of the results found in the main specification, and to investigate whether an effect exists in

Figure 17: Outcome: Work in Office



the cases where a differential trend between exposed and non-exposed hospitalists may be driving results.

I place restrictions on how large the violations in parallel trends must be in order to change my findings. I plot the effect of EHR on each outcome in the first year after exposure under various violations in parallel trends in Figure 24. $M=0$ indicates a linear violation in the parallel trends assumption. The interpretation of violations in parallel trends depends on the magnitude of the confounder. For example, if a one standard deviation in manager productivity hypothetically led to an increase of one extra patient per week (52 per year), then $M = 12$ corresponds to allowing the slope of the differential trend to change by .23 standard deviations of manager productivity. Since we are considering Medicare patients alone, this is a relatively large magnitude. In other words, it would take a large endogenous shock to productivity to explain the effect found on patient count. A similar conclusion can be drawn for both office outcomes, where various relaxations of the parallel trends assumption can be made before ruling out the effect of EHR exposure. When considering exiting clinical work, even the assumption of linear violations in parallel trends rules cannot rule out no effect of EHR exposure on exiting. However, this is not surprising due to the rarity of exiting in the sample. Further, when we consider the whole story of EHR exposure on exit, work setting, and patient count, the magnitude of a confounder would have to be large to explain all of these results together.

Figure 18: Outcome: Fraction of Patients in Office

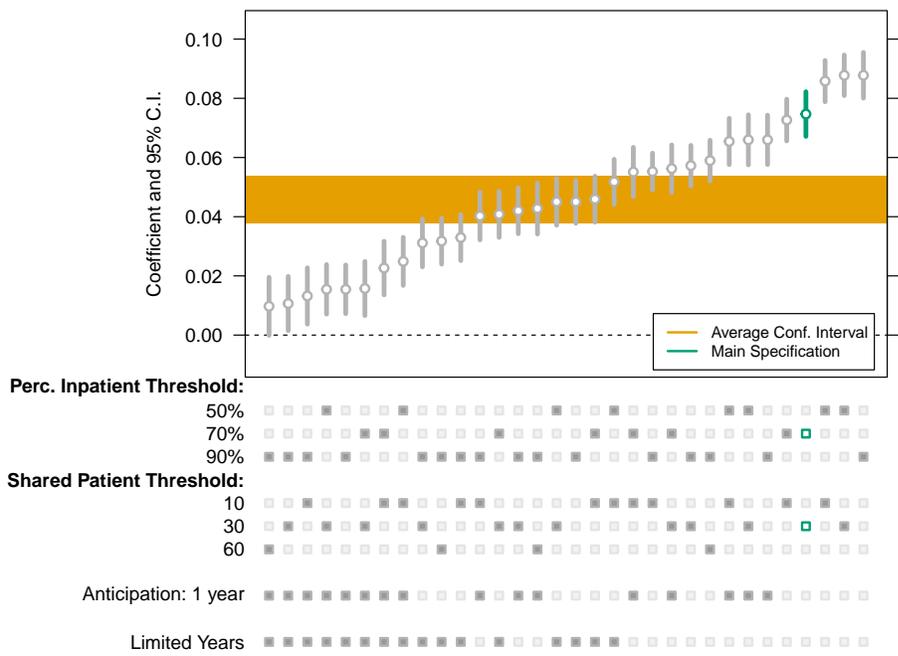
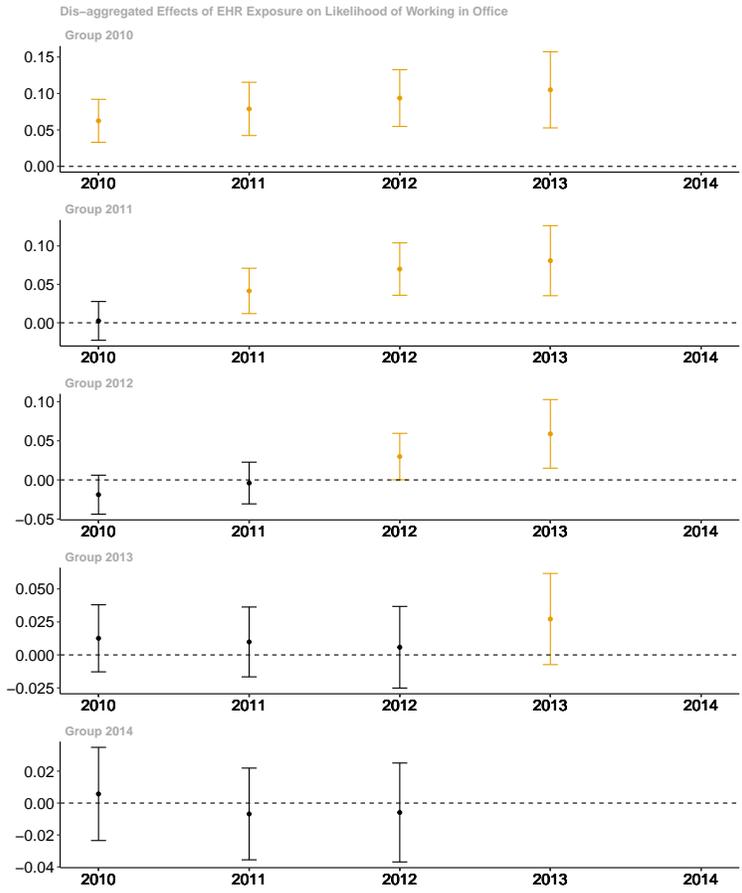
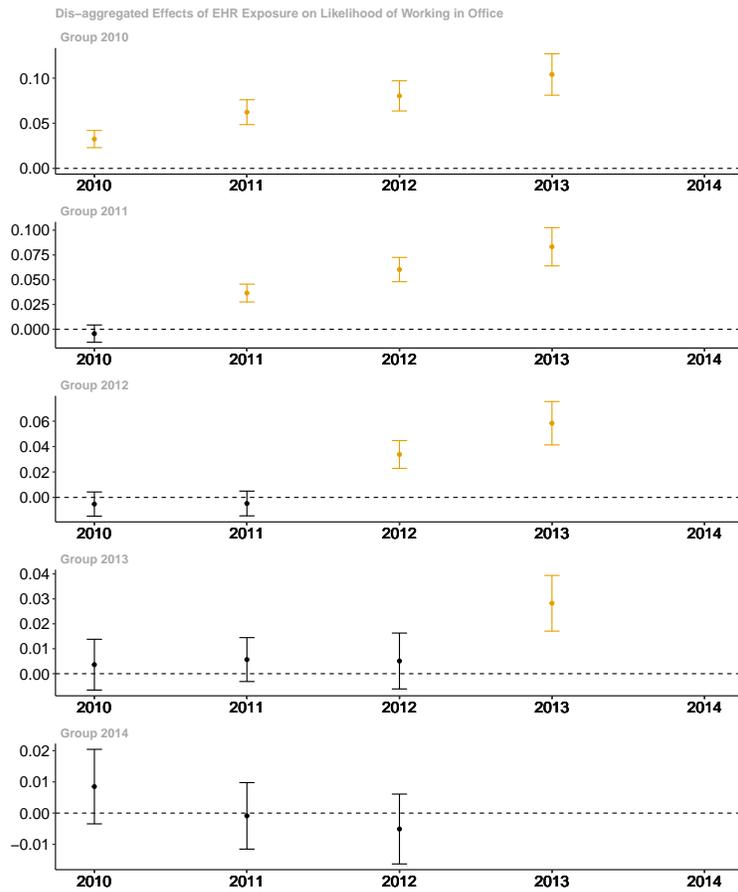


Figure 19: Effect of EHR Exposure on Working in Office by Group



Notes: Plot shows average group time treatment effects for each treatment group.

Figure 20: Effect of EHR Exposure on Frac Patients in Office by Group



Notes: Plot shows average group time treatment effects for each treatment group.

Figure 21: Outcome: Patient Count

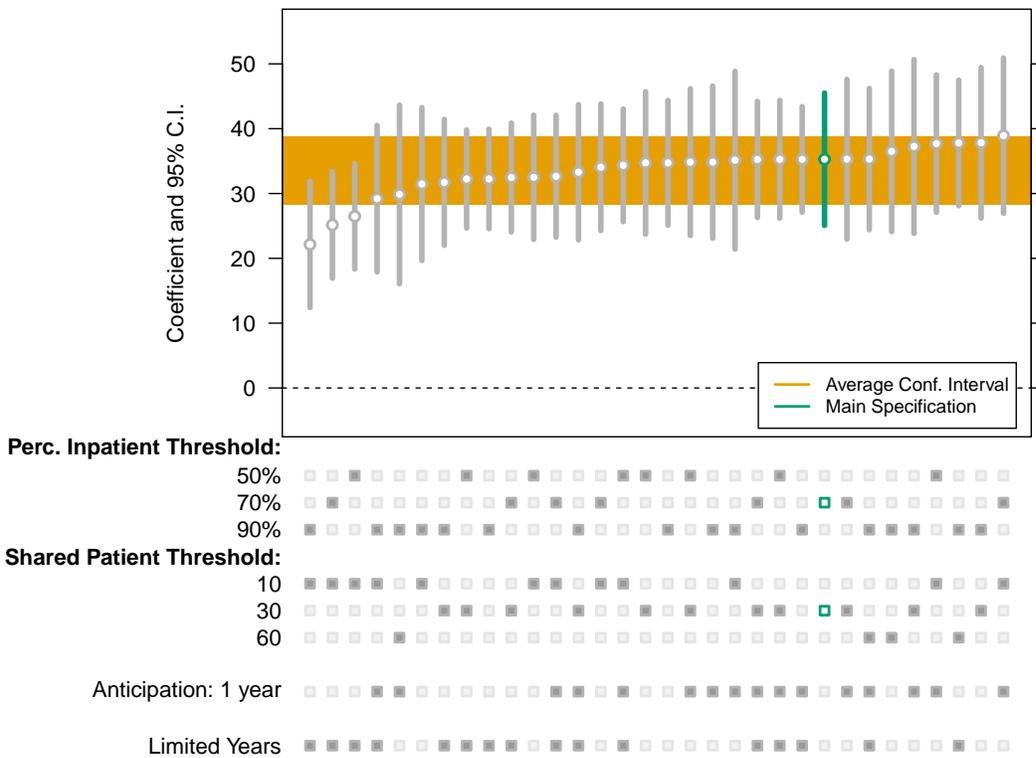


Figure 22: Outcome: Claims per Patient

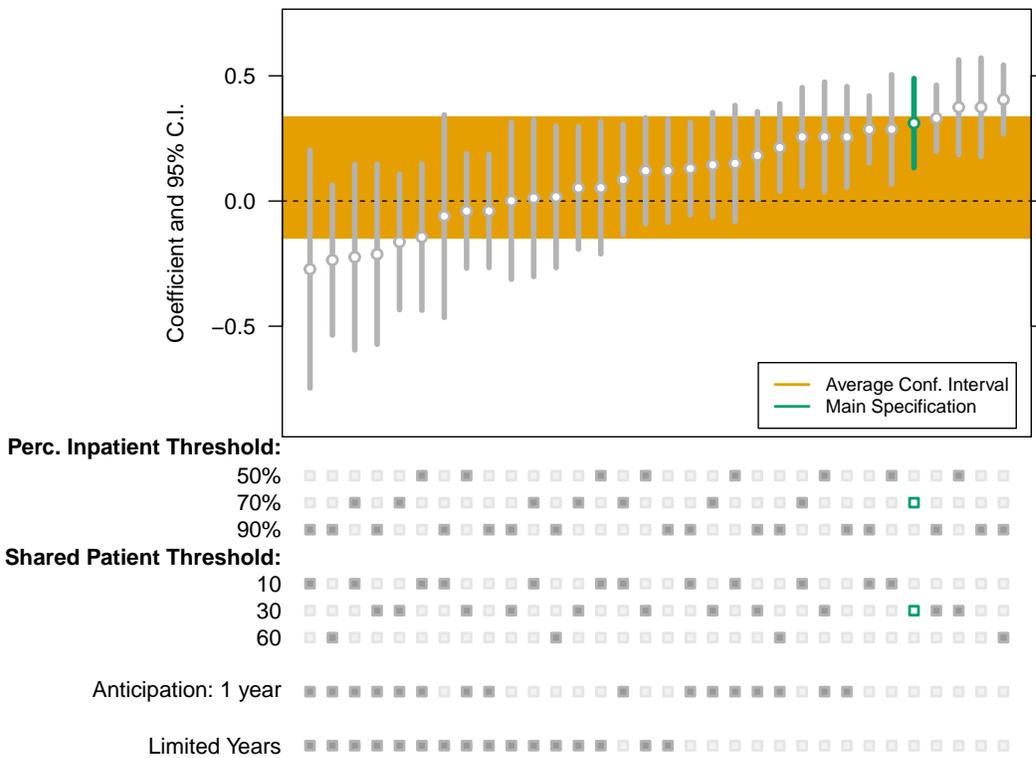
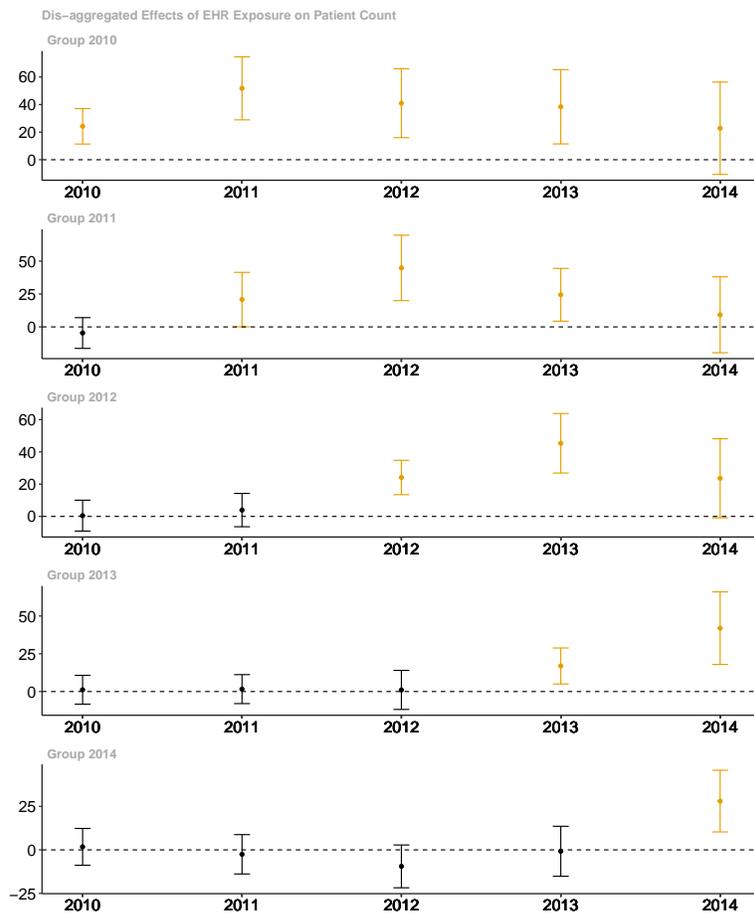


Figure 23: Effect of EHR Exposure on Patient Count by Group



Notes: Plot shows average group time treatment effects for each treatment group.

Figure 24: Effects under Parallel Trends Violations

