The Impacts of AI Integration on Efficiency and Employment Patterns in American Hospitals*

Daniel R. Arnold† Jonathan Cantor‡ Christopher M. Whaley§

March 25, 2026

Abstract

Since the launch of OpenAI's ChatGPT in November 2022, companies across the global economy have rapidly adopted generative artificial intelligence (AI) tools. Adoption of AI technologies is believed to improve business efficiencies and operations. Because efficiencies are expected to result from tasks traditionally performed by humans, AI adoption may reduce employment. Both aspects of AI adoption's impact lack substantial existing empirical evidence, largely due to a lack of comprehensive national data on AI adoption and employment patterns. Studying the impact of AI adoption on employment in the health care sector is especially important, since administrative tasks within the sector appear particularly prone to automation and disruption by AI adoption. However, it remains unclear whether reductions in human labor also affect the quality and delivery of health care. Combining data from 2021-2024 on AI adoption across all U.S. hospitals with measures of hospital performance, we find that AI adoption led to a 3.3% increase in net patient revenue, a 1.2% decrease in administrative expenses, and a 6.6% increase in patient volume. We further find that the observed improvements in clinical quality, taken in aggregate, imply that AI adoption in hospitals reduced mortality by roughly 15,000 deaths nationwide.

*Supported by NIA R01AG073286. We thank Ben Ukert, Andy Ryan, and Seminar participants at AHEC and Brown for helpful comments.

†Brown University, daniel_arnold@brown.edu

‡RAND, jcantor@rand.org

§Brown University, christopher_whaley@brown.edu

1 Introduction

The use of generative artificial intelligence (GenAI) has rapidly accelerated across the global economy (Brynjolfsson, Li and Raymond, 2025; Singla et al., 2025; Humlum and Vestergaard, 2025*a*). This acceleration has been spurred by the release of the Generative Pre-trained Transformer (ChatGPT) by OpenAI in November 2022 (Jackson, 2023). Since 2021, an estimated $1.1 trillion has been invested in AI technologies (Stanford University Human-Centered Artificial Intelligence, 2025). Across industries, the potential impacts of AI are multifaceted and uncertain (Brynjolfsson, Chandar and Chen, 2025; Buenaventura et al., 2025; Humlum and Vestergaard, 2025*b*). The adoption of AI technologies can improve firm efficiencies and operations (Babina et al., 2024). At the same time, these efficiencies often are expected to arise from the reduction of or replacement of human-performed tasks, which may reduce human employment (Acemoglu, Autor and Johnson, 2023; Autor, 2024; Frank et al., 2019; Eloundou et al., 2024). Thus, a fundamental issue is whether the adoption of AI technologies will lead to the substitution of human labor and expertise or contribute to their augmentation (Hui, Reshef and Zhou, 2024). Although AI adoption has been underway for more than a decade, the impacts of this adoption on both firm performance and labor markets remain poorly understood (Hampole et al., 2025).

A particularly important industry to study the impacts of AI adoption is the health care sector (Wachter and Brynjolfsson, 2024; Chen and Zink, 2025; Zink et al., 2023; Nguyen et al., 2025). Health care firms have many tasks where AI technologies are expected to improve performance (e.g., using AI to interpret radiology images). However, automation within the sector is expected to differ based on the skills and roles of workers (Dranove and Garthwaite, 2024). The health care industry also has a high proportion of administrative tasks (e.g., medical billing and clinical note taking) that are currently manually performed and can be automated or improved using new technologies (Sahni et al., 2024). One recent report suggested that hospitals spend approximately 56% of their total operating revenue on labor (Deloitte, 2024).

As a result, there is considerable potential that GenAI can both improve performance and reduce administrative burden in the U.S. health care system (Howell, Corrado and DeSalvo, 2024; Shuaib, 2024). In a recent survey by the American Medical Association (AMA), more than half of physicians (57%) reported that reducing administrative burdens through automation was the biggest area of opportunity for AI (American Medical Association, 2025). Other areas of opportunity include improving patient care outcomes by increasing the speed and accuracy of diagnoses (Jiang, Edwards and Newstead, 2021; Khalifa and Albadawy, 2024) and patient engagement and adherence (Babel et al., 2021; Sekandi et al., 2023). Yet, previous research has primarily focused on rates of self-reported AI adoption within the health care sector (Nguyen et al., 2025; Nong et al., 2025; Bin Abdul Baten, 2024; Goldfarb, He and Teodoridis, 2025; Hwang et al., 2026). Despite the potential of GenAI, there is currently little evidence on the impact of AI adoption and measures of system performance.

In this paper, we examine the multi-dimensional impacts of AI adoption on one of the most important industries in the US health care–hospitals. We leverage national data from several sources that include detailed measures of AI adoption, as well as granular measures of firm (e.g., hospital) performance, including patient throughput, staffing, and capacity. We also measure hospital quality using national data on hospital quality metrics. A key challenge with examining this question is the potential endogeneity between AI adoption and firm performance. We leverage a synthetic difference-in-differences approach that combines both standard difference-in-differences with synthetic controls approaches (Arkhangelsky et al., 2021). As an additional test, we examine the robustness of our results to using more standard event study approaches (both OLS two-way fixed effects and Sun-Abraham estimators) and find similar results.

We find several key impacts of AI adoption. First, we find a 3 percent increase in hospital net patient revenue, which is driven by a nearly 7 percent increase in patient volume. We also find evidence that AI adoption both compliments high-performance human labor, but also substitutes for more routine human labor. Overall employee full-time equivalents and

total labor hours increase by approximately 1 and 3 percent, respectively. Yet, administrative hours decrease by nearly 3 percent. Finally, we find meaningful reductions in clinical quality–heart failure, pneumonia, COPD, and stroke mortality decrease by approximately 0.2 percentage points. In aggregate, our findings imply adoption of AI tools by hospitals leads to a $42 billion increase in patient revenue, but an offsetting $38 billion increase in operating expenses. Across all measures of mortality, we find an approximate 15,000 reduction in deaths.

Within this set of overall results, we also examine sources of heterogeneity. We find AI adoption leads to larger revenue increases for for-profit, as well as system-affiliated hospitals. We also find larger and more consistent increases in patient discharges among patients with either Medicare or private sources of insurance. We find smaller and noisier increases in lower-paying Medicaid patient discharges.

Our study makes several important contributions. First, the emergence of AI has led to high-profile policy discussions on both impacts on firm performance and questions of whether AI technologies will serve as substitutes or complements for human-performed tasks. Yet, despite the importance of these questions, empirical evidence remains limited. Perhaps most closely related to our study is Brynjolfsson, Li and Raymond (2025), which examines the impacts of AI adoption on productivity of approximately 5,000 customer-support workers, but does not examine impacts on broader labor markets or firm performance. This study is also limited to a narrow worker population. Related studies find evidence that AI spurs firm innovation (Babina et al., 2024, 2025) Other work finds firm exposure to ChatGPT led to a 0.4 percent increase in firm public market valuation and a decrease in worker wages. The literature on the immediate impacts of generative AI builds upon studies of earlier technologies. Autor (2024) examine technological innovations since 1940 and find labor market segmentation based on worker skills.

Second, we more narrowly contribute to the nascent literature specific to how AI tools can be used within the health care sector. Much of what is currently known about the

impacts of recent use of AI by hospitals is related to the characteristics that are associated with adoption: system member hospitals and hospitals with higher operating margins are more likely to be AI adopters (Nong et al., 2025; Bin Abdul Baten, 2024; Goldfarb, He and Teodoridis, 2025), while hospitals located in areas with a high Social Deprivation Index, hospitals located in provider shortage areas, critical access hospitals, and rural hospitals are less likely to use AI (Nong et al., 2025; Chen and Yan, 2025). One recent study quantified trends in the quality of care delivered by hospitals that adopted AI and found mixed results (Hwang et al., 2026).

While differences in hospital adoption of AI technologies are well documented, the impacts of hospital AI adoption on hospital labor and performance are not known. If hospital adoption of AI materially changes hospital finances or operations, then underlying resource disparities across providers may be exacerbated. Further, if AI tools, in addition to impacting hospital operations, also have impacts on patient outcomes or quality of care, then disparities in health outcomes are also likely to be exacerbated (Nong et al., 2025; Nong, Maurer and Dwivedi, 2025).

As AI technologies become prevalent, understanding the impacts of AI adoption on both firm (i.e., hospital) and consumer (i.e., patients) outcomes is critical. AI can do some human jobs, but it is not clear if it complements or substitutes human labor. It is also not clear if new "labor" from AI allows firms to expand capacity or increase throughput. Finally, it is not clear if AI efficiencies are to the firm or if they impact patient care. To inform policymakers, health care administrators, and health care providers, this paper uses a novel synthetic difference-in-differences quasi-experimental approach to examine the impacts of hospital adoption of AI tools on hospital performance.

2 Data

We used data from the American Hospital Association (AHA) Annual Survey to classify hospitals as AI adopters or non-adopters. The AHA Annual Survey began asking questions related to AI in 2021. We used 2021 to 2024 AHA Survey data to track AI adoption by hospitals. The 2024 Annual Survey included hospital characteristics for all 6,100 hospitals in the United States, of which 5,121 were community hospitals (i.e., nonfederal, short-term general, and other specialty hospitals) (American Hospital Association, 2026). For this study, we only considered survey responses from the 4,214 short-term general hospitals within the larger group of community hospitals as they accounted for an overwhelming majority (95%) of community hospital admissions and are more easily comparable (in terms of services offered) than specialty and long-term hospitals.

The survey questions related to AI changed for the 2024 survey and contained more detail than the AI question in the 2021-2023 surveys which simply asked "Does your hospital use artificial intelligence (AI) or machine learning in the following? (Check all that apply): 1) predicting staffing needs, 2) predicting patient demand, 3) staff scheduling, 4) automating routine tasks, 5) optimizing administrative and clinical workflows, 6) none of the above."

In 2024 the survey asked about clinical vs. operational use of AI and whether hospitals were not implementing, exploring, piloting/testing, expanding, or fully integrated (see Table A1 for the survey questions).

Among the short-term general hospitals in the United States in 2024, 2,406 (57%) responded to this question. The response rate was similar in 2023 (53%) and 2022 (56%), but lower in 2021 (20%). Among all XX short-term general hospitals that appear in the data between 2021 and 2024, XX (XX%) answered the question in at least one year of the survey.

We considered a hospital as an AI adopter if it reported it used AI at all (i.e., checked any box 1-5 in the question during FY2021-FY2023). Hospitals that never answered the question or answered the question by checking box 6 (none of the above) were considered to be non-adopters. In 2024, we considered a hospital an adopter if it said AI was "expanding" or "fully

integrated". If a hospital had said it used AI during FY2021-FY2023, but then showed up as saying anything less than expanding and fully integrated in 2024 (e.g., exploring, piloting/testing), then we removed the hospital from the adopter group under the notion that we had likely misclassified the hospital as an adopter in earlier years when the survey included less detail on AI use.

Healthcare Cost Report Information System (HCRIS) data were the source of the hospital performance measures we analyzed (RAND Corporation, 2018). The HCRIS data is a comprehensive collection of financial and statistical information that Medicare-certified institutional health care providers are required to submit annually to the Centers for Medicare & Medicaid Services (CMS). For hospitals, the data provide a detailed breakdown of a facility's operations, including its costs, charges, utilization, and finances. We extracted 2017-2024 data for eight primary outcomes: net patient revenue, total expenses, expenses for clinical activities, expenses for administrative activities, discharge equivalents (i.e., volume), employees on payroll (full-time equivalents), total paid hours, and administrative and general paid hours (Prager and Schmitt, 2021; Schmitt, 2017).

To measure care quality, we used 2017-2024 quality of care data from CMS's Hospital Compare dataset. We assessed the association between AI adoption and seven quality outcome measures: six mortality measures – the 30-day mortality rate after being admitted to the hospital for heart attack, heart failure, pneumonia, Chronic Obstructive Pulmonary Disease (COPD), stroke, or Coronary Artery Bypass Graft (CABG) surgery, respectively – and a hospital's all-cause 30-day readmission rate. We additionally extracted five patient experience measures – communication with nurses, communication with doctors, receiving help, rating the hospital a 9 or 10, and recommending the hospital. The mortality and readmission measures are commonly used in the literature (Oseran et al., 2023; Werner and Bradlow, 2006), as are the patient communication and hospital recommendation measures (Jha et al., 2008).

3 Empirical Approach

A key challenge when assessing the impacts of AI on hospital performance is that hospitals that adopt AI likely differ from those that do not. For example, adopters may have more financial resources or operate in more friendly payer environments. As a solution, we used the Synthetic Difference-in-Differences (SDID) method to estimate the association between AI adoption and changes in our outcome measures (Arkhangelsky et al., 2021).

The goal of SDID is to estimate the Average Treatment Effect on the Treated (ATT) by constructing a credible counterfactual for what would have happened to the treated units in the absence of treatment. From synthetic control (SC), SDID borrows the idea of reweighting untreated (control) units to create a "synthetic" control group. This synthetic control is constructed so that its outcome trends in the pre-treatment period closely match the trends of the treated group. SDID incorporates the core logic of difference-in-differences (DID) by allowing for time-invariant differences between the treated and control groups (i.e., unit fixed effects) meaning the method focuses on matching pre-treatment trends rather than requiring a perfect match in pre-treatment outcome levels, relaxing a key assumption of the standard SC method.

A key innovation of SDID is the introduction of time weights. In addition to weighting control units, SDID also weights the pre-treatment time periods. It gives more importance to periods that are more predictive of outcomes in the post-treatment period, which can improve the estimator's precision and robustness.

The core estimating equation is:

$$(\hat{\tau}^{sdid}, \hat{\mu}, \hat{\alpha}, \hat{\beta}) = \underset{\tau,\mu,\alpha,\beta}{\arg\min} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} (Y_{it} - \mu - \alpha_i - \beta_t - W_{it}\tau)^2 \hat{\omega}_i^{sdid} \hat{\lambda}_t^{sdid} \right\},$$

where Y_{it} is the outcome for unit i at time t, W_{it} indicates treatment, and α_i and β_t are unit and time fixed effects, respectively. This hybrid approach makes SDID robust in settings where traditional DID may fail due to non-parallel trends and where traditional SC may fail

because the treated unit's outcomes are outside the range of the control units.

In this study, we were faced with both non-parallel trends and differences in pre-treatment outcome levels. The trends were frequently different across treated and control groups and there were sizable level differences given treated hospitals are generally larger.

Our analysis utilizes a balanced panel of 4,093 U.S. general acute care hospitals spanning the years 2017 through 2024. The treated group consists of 1,511 hospitals that adopted clinical or operational AI tools between 2021 and 2024, reflecting a staggered adoption timing. The control group comprises the remaining 2,582 hospitals that did not adopt AI during this period.

Estimation proceeded in two primary phases. First, we applied the SDID algorithm separately for each adoption cohort (e.g., hospitals strictly adopting in 2021). For each cohort, the pre-treatment period was defined as the years from 2017 up to the year prior to adoption, with the control group consisting of all 2,582 non-adopting hospitals. Unit weights ($\hat{\omega}$) were calculated to construct a synthetic control from the non-adopting hospitals, ensuring that the average pre-adoption outcome trends were parallel to those of the specific adoption cohort. Additionally, time weights ($\hat{\lambda}$) were calculated for the pre-adoption years to improve the precision of the counterfactual prediction. Using these weights, we estimated a cohort-specific treatment effect via a weighted difference-in-differences regression. Second, to produce a single, overall ATT, we averaged the individual treatment effects from each cohort, weighted by the size of the cohort.

Building on this foundation, the SDID framework can be extended to an event-study analysis to understand the dynamic, period-by-period effects of a treatment. This extension disaggregates the overall ATT into a series of estimates, showing how the treatment effect evolves over time after implementation.

The core of the event-study extension is the dynamic cohort-specific estimator ($\hat{\tau}^{sdid}_{a,l}$), which measures the treatment effect for a specific adoption cohort (a) at a specific time lag (l) after treatment begins. Instead of averaging over all post-treatment periods, this estimator

compares the outcome difference between a treated cohort and its synthetic control at a single point in time $(a-1+l)$ to the average difference observed across the entire pre-treatment period. The estimator for the effect l periods after treatment for cohort a is defined as:

$$\hat{\tau}_{a,l}^{sdid} = \left(\frac{1}{N_{tr}} \sum_{i=1}^{N_{tr}} Y_{i,a-1+l} - \sum_{j=1}^{N_{co}} \hat{\omega}_j Y_{j,a-1+l} \right) - \text{Pre-treatment difference}$$

By construction, the standard cohort-specific SDID estimate ($\hat{\tau}_a^{sdid}$) is simply the average of these dynamic estimators over all post-treatment periods for that cohort.

In a staggered adoption setting where different cohorts receive treatment at different times, it is useful to have a single estimate for the effect at a specific lag l (e.g., the average effect two years post-treatment). This is achieved by calculating the aggregated dynamic estimator ($\hat{\tau}_l^{sdid}$), which is a weighted average of the cohort-specific dynamic effects ($\hat{\tau}_{a,l}^{sdid}$) across all cohorts observed l periods after their treatment began. The weights are based on the relative number of treated units in each cohort, giving more influence to larger cohorts. The formula for this aggregated dynamic estimator is:

$$\hat{\tau}_l^{sdid} = \sum_{a \in A_l} \frac{N_{tr}^a}{N_{tr}^l} \hat{\tau}_{a,l}^{sdid},$$

where N_{tr}^l is the total number of treated units across all cohorts observed at lag l. This framework is internally consistent, as the overall ATT from the standard SDID procedure can be recovered from these dynamic estimates. The overall ATT is a weighted average of the aggregated dynamic estimators ($\hat{\tau}_l^{sdid}$) with weights proportional to the total number of units for which the effect at lag l can be calculated. This relationship is shown by the equation:

$$ATT = \frac{1}{T_{post}} \sum_{l=1}^{T_{tr}} N_{tr}^l \hat{\tau}_l^{sdid},$$

which demonstrates how the overall average effect is composed of the period-by-period dy-

namic effects. For completeness, we also estimate traditional two-way fixed effects (TWFE) and Sun and Abraham (2021) versions of our event studies and present the results in the appendix.

We conducted several additional analyses to analyze where AI adoption is having the largest impact. Specifically, we tested whether the impact of AI adoption differed by hospital profit status (for-profit vs. nonprofit), system membership, and critical access status. To be designated a critical access hospital, a hospital must have less than 25 beds and be geographically separate from other hospitals. Critical access hospitals have unique financing models and face financial pressures that have contributed to closures or consolidation with health systems (Whaley, Bartlett and Bai, 2024). We hypothesize that to date for-profit hospitals, hospitals that are part of systems (as opposed to independent hospitals), and non-critical access hospitals have benefited more from AI adoption. We also analyzed how AI adoption impacts patient volume by service type (i.e., inpatient or outpatient) and payer (i.e., commercial, Medicare, or Medicaid).

4 Results

4.1 Main Results

We identified 1,511 hospitals that adopted AI over the 2021-2024 period. 624 hospitals first reported using AI in 2021, while 165, 100, and 622 hospitals adopted in 2022, 2023, and 2024, respectively. Table A1 in the appendix compares characteristics of the AI adopter and non-adopter hospitals at baseline in 2017. AI adopter hospitals were larger (221 beds vs. 114 beds), more likely to be part of a system (0.87 vs. 0.61), less likely to be critical access hospitals (0.19 vs. 0.39), more often located in urban areas (urban continuum code 2.7 vs. 4.1), less likely to be for-profit (0.07 vs. 0.22), more often a major teaching hospital (0.13 vs. 0.05), and had higher casemix indices (1.68 vs. 1.56) than non-adopter hospitals.

Hospital finances and volume. AI adoption was associated with a 3.3% (95% confidence interval [CI], 2.4% to 4.1%) increase in net patient revenue, a 2.9% (95% CI, 2.1% to 3.7%) increase in total expenses, a 2.8% (95% CI, 2.0% to 3.7%) increase in clinical expenses, a 1.2% (95% CI, -2.5% to 0.1%) decrease in administrative expenses, and a 6.6% (95% CI, 5.2% to 8.0%) increase in discharge equivalents (Figure 1). We used log values for the dependent variables and then transformed them to percent changes for each of these five results.

The event study plots for net patient revenue and administrative expenses are shown as Figures 2 and 3, respectively. Figure 2 shows that AI adoption immediately increased net patient revenue (1.4%; 95% CI, 0.8% to 2.1%) and the effect continued to show signs of growing three years post-adoption when the magnitude reached 6.6% (95% CI, 4.6% to 8.6%). 6.6% of the adopting hospital baseline mean net patient revenue ($339 million) equates to $22 million in additional revenue. Figure 3 shows AI adoption reduced administrative expenses by 1.9% after two years (95% CI, -2.1% to 0.1%) with estimates becoming noisier after two years post. 1.9% of the adopting hospital baseline mean administrative expenses ($60 million) is a $1.1 million reduction in administrative expenses. TWFE and Sun-Abraham versions of 2 and 3 (shown in the appendix as Figures A10 and A11) follow the same pattern as their corresponding SDID versions, but with clear signs of pre-trends that necessitated our SDID approach. Event study plots for the other finance and volume outcomes are shown in Figure A2 in the appendix.

Labor. AI adoption was associated with a 1.3% increase in employee FTEs (95% CI, 0.2 to 2.4), a 3.2% increase in total paid hours (95% CI, 2.1 to 4.3), and a not statistically significant 1.9% (95% CI, -4.4 to 0.6) decrease in administrative paid hours (Figure 1) (though the coefficient four years post (-6.6%) is statistically significant). Event study plots for the labor outcomes are shown in Figure A3 in the appendix.

Quality. AI adoption was associated with a 0.22 percentage point (pp) (95% CI, -0.31 to -0.14) decrease in the heart failure mortality rate (Figure 4). The mean heart failure

mortality rate at baseline for AI adopters was 11.6% so this corresponds to a 1.9% relative decrease. AI adoption was also associated with a 0.41 pp (95% CI, -0.52 to -0.30) decrease in the pneumonia mortality rate, a 0.21 pp (95% CI, -0.28 to -0.14) decrease in the COPD mortality rate, and a 0.22 pp (95% CI, -0.30 to -0.14) decrease in the stroke mortality rate, which correspond to 2.5%, 2.4%, and 1.6% relative decreases, respectively. The other three quality outcome measures we analyzed were not associated with AI adoption. Three of the five quality patient experience measures we analyzed likewise were not associated with AI adoption, with nurse communication and "receiving help" being the exceptions. AI adoption was associated with a 0.44 pp (95% CI, -0.73 to -0.15) decrease in the percentage of patients who said their nurses always communicated well. At baseline, 66% of patients at AI adopting hospitals reported their nurses always communicated well, so this effect corresponds to a less than 1% relative decrease. AI adoption was associated with a 0.21 pp (95% CI, -0.89 to -0.13) decrease in the percentage of patients who said they received help when needed. At baseline, 79% of patients at AI adopting hospitals reported receiving help when needed, so this effect corresponds to a less than 1% relative decrease. Event study plots for quality outcomes and patient experience associations are shown in Figures A4 and A5, respectively.

Table 1 places the previously discussed results in an aggregate context. The 1,511 adopting hospitals were treated for 2.5 years on average of the possible four years of treatment (2021-2024). We multiplied 1,511 by 2.5 and the appropriate means and coefficient estimates from Table A1 and Figure 1 to produce the aggregate results in Table 1. For instance, for the aggregate net patient revenue effect we multiplied 1,511 by 2.5 and $339 million (mean net patient revenue at baseline for AI adopters in Table A1) and 0.033 (the percent change in net patient revenue from Figure 1). Using this procedure, our estimated aggregate effects are $42.0 billion (95% CI, $31.2 billion to $52.9 billion) in net patient revenue, $38.0 billion (95% CI, $27.7 billion to $48.4 billion) in total expenses, -$2.7 billion (95% CI, -$5.6 billion to $0.2 billion) in administrative expenses, 4.7 million (95% CI, 3.7 million to 5.7 million) discharge equivalents, and 61,000 (95% CI, 40,000 to 82,000) additional employees. Addi-

tionally, we estimated over 15,000 fewer in-hospital deaths (95% CI, 10,803 to 20,661) across the four mortality measures that improved following AI adoption. Additional details on how we calculated aggregate effects are available in the appendix.

4.2 Additional Results

We additionally find AI adoption leads to larger increases in revenues among for-profit, non-critical access, and system hospitals, but subgroup estimates were not statistically different from one another, largely because of reduced sample size. We found little heterogeneity across the same three dimensions in terms of effects on administrative expenses (Figures A6 and A7 in the appendix).

We also analyzed whether the volume increases we identified were contained to a particular service type (e.g., inpatient) or payer (e.g., commercial) or applied more broadly. Figures A8 and A9 in the appendix show volume increased across the board – by type (inpatient, outpatient) and payer (commercial, Medicare, and Medicaid). Notably, hospital beds did not increase at adopting hospitals until three years post-adoption and when bed expansion did occur it was at a much smaller level (2.0%; 95% CI, 0.3% to 3.9%) than the increase in patient volume.

Table A2 in the appendix gives a sense of what types of clinical and operational AI hospitals are adopting. Among the over 2,000 hospital respondents to the 2024 AI survey questions. By summing the "expanding" or "fully integrated" columns, the data makes it clear there is a preference for specific clinical applications. Notably, Clinical Decision Support Tools (43%), Predictive Analytics for Patient Care (39%), and AI-Assisted Diagnostics (37%) demonstrate substantially higher adoption rates compared to the rest of the surveyed applications. While operational applications generally lag behind, Revenue Cycle Management does show some meaningful traction, with 29% of hospitals reporting it as expanding or fully integrated. Overall, however, the highest rates of deep integration are concentrated within patient-facing and diagnostic tools, suggesting that advanced clinical

AI implementation is currently more common and mature than operational AI use cases in hospitals.

5 Conclusion

The launch of ChatGPT and other AI applications has created the potential to transform the health care industry by improving patient interaction, supporting clinical diagnoses, enhancing health education, providing health advice, promoting wellness, and increasing efficiencies (Fui-Hoon Nah et al., 2023). This study shows that the adoption of AI by hospitals is associated with increases in revenue, clinical expenses, and volume and decreases in administrative costs. Additionally, we show evidence that hospital AI adoption is associated with 15,000 fewer in-hospital deaths. Finally, this paper finds that AI adoption increases patient volume across service and payer type and that this additional volume is largely a result of improved efficiency rather than physical expansion (e.g., more hospital beds).

There are clear differences in the characteristics of hospitals that have adopted AI and those that have not. AI adopting hospitals are larger, more urban, more likely to be a major teaching hospital, and treat more complex patient populations, which is consistent with previous studies (Nong et al., 2025; Bin Abdul Baten, 2024; Goldfarb, He and Teodoridis, 2025; Poon et al., 2025). They are also much less likely to be critical access hospitals than non-adopters. Organizations with greater resources are more capable of financing AI adoption and hiring staff with AI development expertise (Nong, Maurer and Dwivedi, 2025; Burns et al., 2025). These differences in AI adoption raises the question of whether smaller, more rural (and presumably financially less stable) hospitals have access to the benefits of AI and what the future looks like for these hospitals, their providers, and their patients if they are unable to access these benefits.

This study's findings should be interpreted while cognizant of its several limitations. First, a number of the hospitals we labeled as non-adopters could actually be adopters who

simply chose not to answer the AHA's survey question on the topic. In this case, they could be receiving weight in our synthetic control and would in turn bias our results toward the null. Second, the SDID method used in this study hinges on the assumption that the constructed synthetic control is a valid counterfactual, that is, it would have continued on the same trend as the treated group absent the treatment.

Initial findings suggest that the integration of AI into hospital administrative workflows holds substantial promise for mitigating systemic operational burdens. Emerging AI applications in areas such as predictive scheduling algorithms, Natural Language Processing for automated clinical documentation and billing, and streamlined revenue cycle management are showing a measurable impact on efficiency. A direct corollary of these gains is the potential for increased patient throughput, which logically translates to enhanced revenue streams. However, while these preliminary results are encouraging, they necessitate a deeper, more systematic inquiry. Future research must be directed toward elucidating the specific causal mechanisms by which AI interventions generate these improvements. A critical challenge moving forward will be developing scalable and equitable implementation strategies to ensure these advancements are not confined to a few large medical centers but can be disseminated throughout the diverse landscape of the American health care system, thereby maximizing their public health impact.

References

Acemoglu, Daron, David Autor, and Simon Johnson. 2023. "Can we have pro-worker AI." Centre for Economic Policy Research. https://cepr.org/system/files/publication-files/191183-policy_insight_123_can_we_have_pro_worker_ai_choosing_a_path_of_machines_in_service_of_minds.pdf.

American Hospital Association. 2026. "Fast Facts on U.S. Hospitals, 2026." American Hospital Association. https://www.aha.org/statistics/fast-facts-us-hospitals.

American Medical Association. 2025. "Physician enthusiasm grows for health care AI." American Medical Association. https://www.ama-assn.org/press-center/ama-press-releases/ama-physician-enthusiasm-grows-health-care-ai.

Arkhangelsky, Dmitry, Susan Athey, David A Hirshberg, Guido W Imbens, and Stefan Wager. 2021. "Synthetic difference-in-differences." *American Economic Review*, 111(12): 4088–4118.

Autor, David. 2024. "Applying AI to rebuild middle class jobs." National Bureau of Economic Research. Working paper 32140.

Babel, Aditi, Richi Taneja, Franco Mondello Malvestiti, Alessandro Monaco, and Shaantanu Donde. 2021. "Artificial intelligence solutions to increase medication adherence in patients with non-communicable diseases." *Frontiers in Digital Health*, 3: 669869.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2024. "Artificial Intelligence, Firm Growth, and Product Innovation." *Journal of Financial Economics*, 151: 103745.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2025. "Artificial Intelligence Makes Firm Operating Performance Less Volatile." *AEA Papers and Proceedings*, 115: 35–39.

Bin Abdul Baten, Redwan. 2024. "How are US hospitals adopting artificial intelligence? Early evidence from 2022." *Health Affairs Scholar*, 2(10): qxae123.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025. "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence." Stanford University. https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. "Generative AI at work." *The Quarterly Journal of Economics*, 140(2): 889–942.

Buenaventura, Maya, Anton Shenk, Anujin Nergui, Sean Mann, Flannery C. Dolan, Karishma V. Patel, Jim Mignano, and Jonathan W. Welburn. 2025. *Artificial Intelligence Adoption and Sectoral Transformation: Implications for Health Care, Financial Services, Climate and Energy, and Transportation.* Santa Monica, CA:RAND Corporation.

Burns, Michael L, Ssu-Ying Chen, Chu-An Tsai, John Vandervest, Balaji Pandian, Paige Nong, David A Hanauer, Andrew Rosenberg, and Jodyn Platt. 2025. "Generative AI costs in large healthcare systems, an example in revenue cycle." *NPJ Digital Medicine*, 8(1): 579.

Chen, Irene Y, and Anna Zink. 2025. "AI tools in human hands: measuring real-world impact in the emergency department."

Chen, Jie, and Alice Shijia Yan. 2025. "Hospital artificial intelligence/machine learning adoption by neighborhood deprivation." *Medical Care*, 63(3): 227–233.

Deloitte. 2024. "How AI can help hospitals strengthen their financial performance and reduce clinician burnout." Deloitte. https://www.deloitte.com/us/en/services/consulting/articles/artificial-intelligence-in-hospitals-financial-performance-clinical-burnout.html.

Dranove, David, and Craig Garthwaite. 2024. "Artificial Intelligence, the Evolution of the Health Care Value Chain, and the Future of the Physician." In *The Economics of Artificial Intelligence: Health Care Challenges. Working Paper Series*. University of Chicago Press.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. "GPTs Are GPTs: Labor Market Impact Potential of LLMs." *Science*, 384(6702): 1306–1308.

Frank, Morgan R, David Autor, James E Bessen, Erik Brynjolfsson, Manuel Cebrian, David J Deming, Maryann Feldman, Matthew Groh, José Lobo, Esteban Moro, et al. 2019. "Toward understanding the impact of artificial intelligence on labor." *Proceedings of the National Academy of Sciences*, 116(14): 6531–6539.

Fui-Hoon Nah, Fiona, Ruilin Zheng, Jingyuan Cai, Keng Siau, and Langtao Chen. 2023. "Generative AI and ChatGPT: Applications, challenges, and AI-human collaboration."

Goldfarb, Avi, Xianda He, and Florenta Teodoridis. 2025. "Patterns of Artificial Intelligence Adoption by Hospitals." Vol. 115, 40–45, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller. 2025. "Artificial Intelligence and the Labor Market." National Bureau of Economic Research Working Paper 33509.

Howell, Michael D, Greg S Corrado, and Karen B DeSalvo. 2024. "Three epochs of artificial intelligence in health care." *Jama*, 331(3): 242–244.

Hui, Xiang, Oren Reshef, and Luofeng Zhou. 2024. "The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market." *Organization Science*, 35(6): 1977–1989.

Humlum, Anders, and Emilie Vestergaard. 2025*a*. "Large Language Models, Small Labor Market Effects." National Bureau of Economic Research Working Paper 33777.

Humlum, Anders, and Emilie Vestergaard. 2025*b*. "The unequal adoption of ChatGPT exacerbates existing inequalities among workers." *Proceedings of the National Academy of Sciences*, 122(1): e2414972121.

Hwang, Yeon-Mi, Madelena Y. Ng, Malvika Pillai, Michelle P. Sahai, and Tina Hernandez-Boussard. 2026. "The Landscape of AI Implementation in US Hospitals." *Nature Health*, 1(1): 99–112.

Jackson, Amber. 2023. "ChatGPT one year on: adoption, transformation, regulation." *AI Magazine*.

Jha, Ashish K, E John Orav, Jie Zheng, and Arnold M Epstein. 2008. "Patients' perception of hospital care in the United States." *New England Journal of Medicine*, 359(18): 1921–1931.

Jiang, Yulei, Alexandra V Edwards, and Gillian M Newstead. 2021. "Artificial intelligence applied to breast MRI for improved diagnosis." *Radiology*, 298(1): 38–46.

Khalifa, Mohamed, and Mona Albadawy. 2024. "AI in diagnostic imaging: revolutionising accuracy and efficiency." *Computer Methods and programs in biomedicine update*, 5: 100146.

Nguyen, Thuy D, Christopher M Whaley, Kosali Simon, Neil Mehta, Hao Yu, Ryan K McBain, Ateev Mehrotra, and Jonathan H Cantor. 2025. "Adoption of Artificial Intelligence in the Health Care Sector." Vol. 6, e255029–e255029, American Medical Association.

Nong, Paige, Eric Maurer, and Roli Dwivedi. 2025. "The urgency of centering safety-net organizations in AI governance." *npj Digital Medicine*, 8(1): 117.

Nong, Paige, Julia Adler-Milstein, Nate C Apathy, A Jay Holmgren, and Jordan Everson. 2025. "Current Use And Evaluation Of Artificial Intelligence And Predictive Models In US Hospitals: Article examines uses and evaluation of artificial intelligence and predictive models in US hospitals." *Health Affairs*, 44(1): 90–98.

Oseran, Andrew S, Rishi K Wadhera, E John Orav, and Jose F Figueroa. 2023. "Effect of Medicare advantage on hospital readmission and mortality Rankings." *Annals of Internal Medicine*, 176(4): 480–488.

Poon, Eric G, Christy Harris Lemak, Juan C Rojas, Janet Guptill, and David Classen. 2025. "Adoption of artificial intelligence in healthcare: survey of health system priorities, successes, and challenges." *Journal of the American Medical Informatics Association*, 32(7): 1093–1100.

Prager, Elena, and Matt Schmitt. 2021. "Employer consolidation and wages: Evidence from hospitals." *American Economic Review*, 111(2): 397–427.

RAND Corporation. 2018. "RAND Hospital Data: Web-Based Tool." RAND Corporation. https://www.rand.org/pubs/tools/TL303.html.

Sahni, Nikhil R., George Stein, Rodney Zemmel, and David Cutler. 2024. "The Potential Impact of Artificial Intelligence on Health Care Spending." In *The Economics of Artificial Intelligence: Health Care Challenges. NBER Chapters*. University of Chicago Press.

Schmitt, Matt. 2017. "Do hospital mergers reduce costs?" *Journal of health economics*, 52: 74–94.

Sekandi, Juliet Nabbuye, Weili Shi, Ronghang Zhu, Patrick Kaggwa, Ernest Mwebaze, and Sheng Li. 2023. "Application of artificial intelligence to the monitoring of medication adherence for tuberculosis treatment in Africa: algorithm development and validation." *JMIR AI*, 2(1): e40167.

Shuaib, Ali. 2024. "Transforming healthcare with AI: Promises, pitfalls, and pathways forward." *International journal of general medicine*, 1765–1771.

Singla, Alex, Alexander Sukhovetsky, Bryce Hall, Lareina Yee, and Michael Chui. 2025. "The state of AI in 2025: agents, innovation, and transformation." McKinsey & Company. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai#/.

Stanford University Human-Centered Artificial Intelligence. 2025. "Artificial Intelligence Index Report 2025." Stanford University. https://hai.stanford.edu/ai-index/2025-ai-index-report.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of econometrics*, 225(2): 175–199.

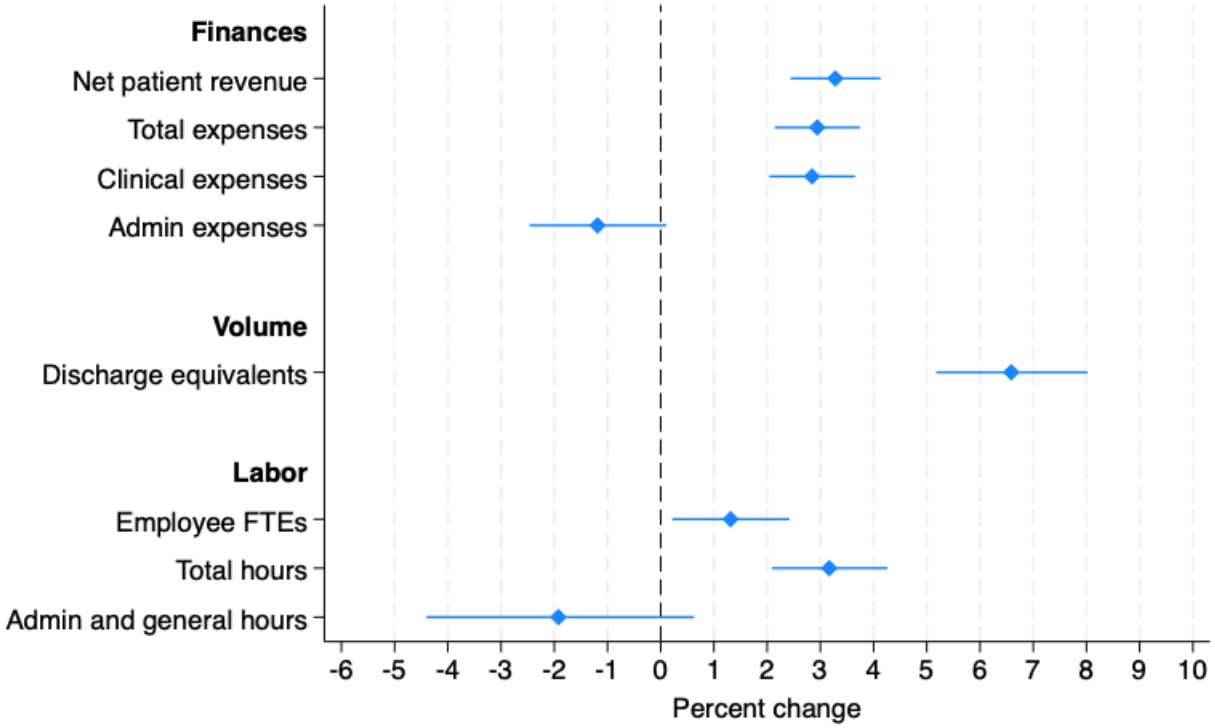
Wachter, Robert M, and Erik Brynjolfsson. 2024. "Will generative artificial intelligence deliver on its promise in health care?" *Jama*, 331(1): 65–69.

Werner, Rachel M, and Eric T Bradlow. 2006. "Relationship between Medicare's hospital compare performance measures and mortality rates." *Jama*, 296(22): 2694–2702.

Whaley, Christopher, Marilyn Bartlett, and Ge Bai. 2024. "Financial performance gaps between critical access hospitals and other acute care hospitals." Vol. 5, e243959–e243959, American Medical Association.

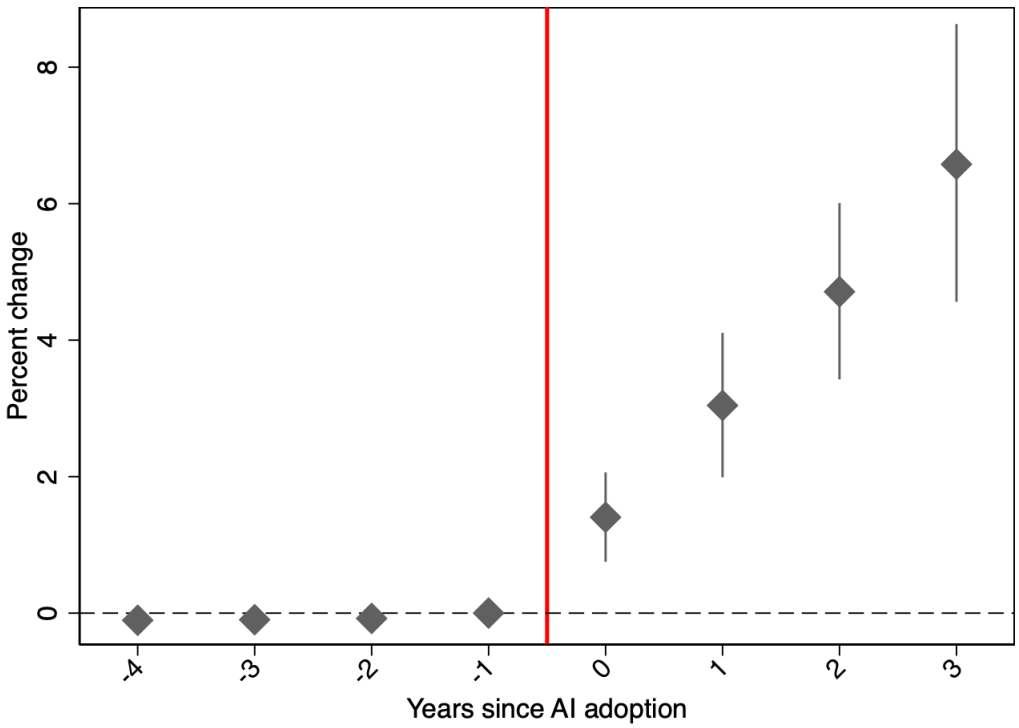
Zink, Anna, Sarah Morriss, Anuj Gangopadhyaya, and Ziad Obermeyer. 2023. "Building Equitable Artificial Intelligence in Health Care." *Urban Institute*.

Figure 1: Summary of AI adoption effects on revenues, expenses, volume, employment, efficiency



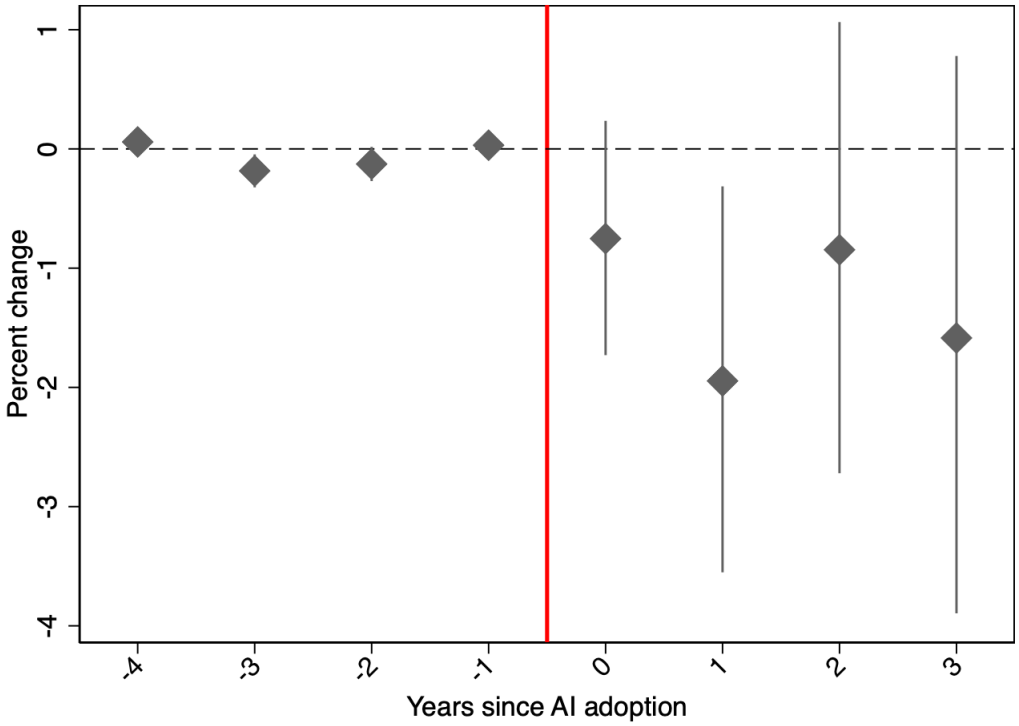
Notes: The figure shows estimates from a synthetic difference-in-differences model where the dependent variable for each model corresponds to the logged version of the variable listed on the vertical axis. Coefficient estimates have been transformed using the formula exp(coef - 1)*100 to create the percentage effects shown in the figure. The regression model included hospital and year fixed and standard errors were clustered at the hospital level. The error bars represent 95% confidence intervals.

Figure 2: Percentage change in net patient revenue after AI adoption



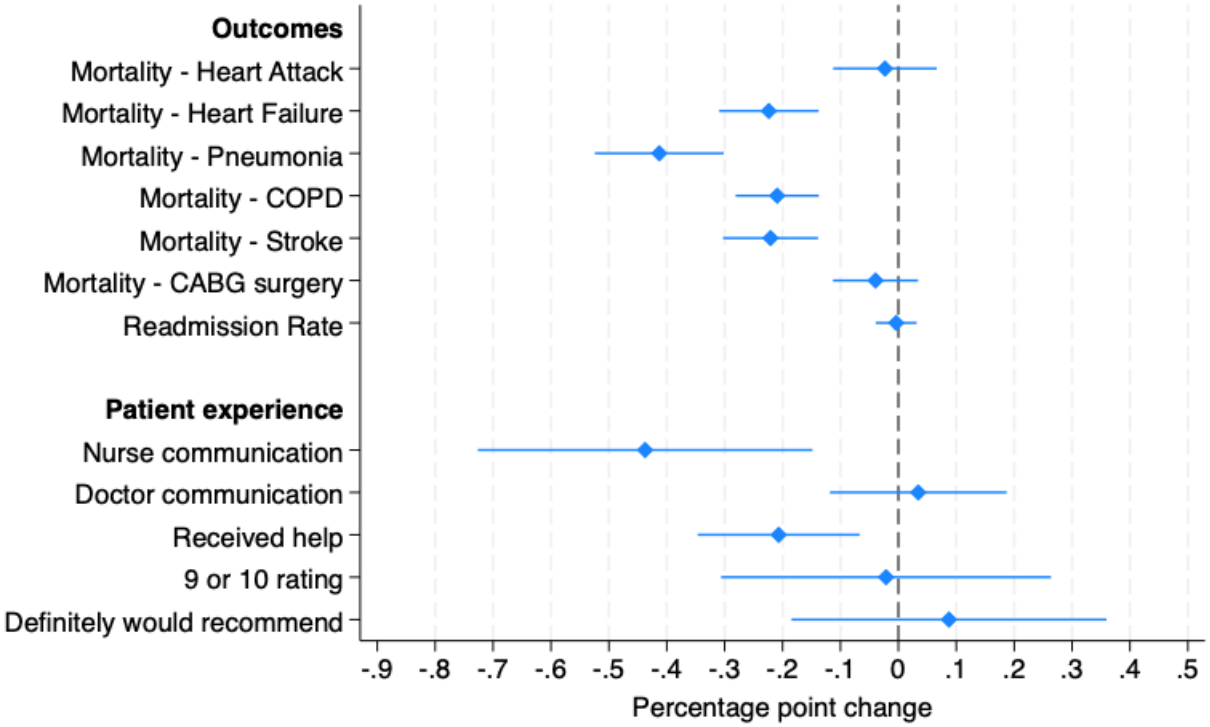
Notes: The figure shows estimates from a synthetic difference-in-differences event study where log(net patient revenue) is the dependent variable. Coefficient estimates have been transformed using the formula exp(coef - 1)*100 to create the percentage effects shown in the figure. The regression model included hospital and year fixed and standard errors were clustered at the hospital level. The error bars represent 95% confidence intervals.

Figure 3: Percentage change in administrative expenses after AI adoption



Notes: The figure shows estimates from a synthetic difference-in-differences event study where log(administrative expenses) is the dependent variable. Coefficient estimates have been transformed using the formula exp(coef - 1)*100 to create the percentage effects shown in the figure. The regression model included hospital and year fixed and standard errors were clustered at the hospital level. The error bars represent 95% confidence intervals.

Figure 4: Summary of AI adoption effects on quality



Notes: The figure shows estimates from a synthetic difference-in-differences model where the variable listed on the vertical axis is the dependent variable for each model. The regression model included hospital and year fixed and standard errors were clustered at the hospital level. The error bars represent 95% confidence intervals.

Table 1: Aggregate effects of hospital AI adoption, 2021-2024

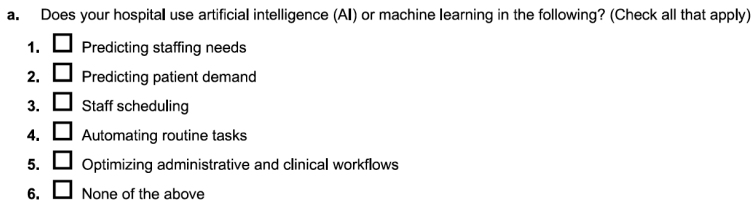
	Estimate (95% CI)
Finances	
Net patient revenue	$42.0 billion ($31.2 to $52.9 billion)
Total expenses (millions of $)	$38.0 billion ($27.7 to $48.4 billion)
Clinical expenses (millions of $)	*$19.1 billion ($13.7 to $24.6 billion)*
Admin. expenses (millions of $)	*-$2.7 billion (-$5.6 to $0.2 billion)*
Volume	
Discharge equivalents	4.7 million (3.7 to 5.7 million)
Labor	
Total hours	490 million (323 to 658 million)
	∼61k employees (40k to 82k employees)*
Quality	
Mortality - heart failure	2,983 fewer in-hospital deaths (1,837 to 4,128)
Mortality - pneumonia	8,809 fewer in-hospital deaths (6,434 to 11,184)
Mortality - COPD	1,953 fewer in-hospital deaths (1,283 to 2,624)
Mortality - stroke	1,987 fewer in-hospital deaths (1,249 to 2,725)

Notes: The table shows aggregated effects calculated by multiplying the SDID estimates by 1,511 (the number of adopting hospitals), 2.5 (the average number of adoption years), and the relevant adopting hospital sample means from Table A1 in the appendix. Full calculation details are included in the appendix. *Assumed to work 2,000 hours (40 hrs/week x 50 weeks) per year over four years.

Appendix

Figure A1: AHA Annual Survey Artificial Intelligence Questions, FY2021 - FY2024

(a) FY2021-FY2023



a. Does your hospital use artificial intelligence (AI) or machine learning in the following? (Check all that apply)

1. ☐ Predicting staffing needs
2. ☐ Predicting patient demand
3. ☐ Staff scheduling
4. ☐ Automating routine tasks
5. ☐ Optimizing administrative and clinical workflows
6. ☐ None of the above

(b) FY2024



a. How would you describe your **hospital's** current level of AI implementation in the following clinical areas?

	(1) Not Implementing	(2) Exploring	(3) Piloting/Testing	(4) Expanding	(5) Fully Integrated	(0) Don't Know
1. AI-assisted diagnostics (including imaging & early detection)	☐	☐	☐	☐	☐	☐
2. Predictive analytics for patient care (including outcomes & deterioration)	☐	☐	☐	☐	☐	☐
3. Clinical decision support tools	☐	☐	☐	☐	☐	☐
4. AI-assisted surgery	☐	☐	☐	☐	☐	☐
5. AI-powered patient communication and education	☐	☐	☐	☐	☐	☐
6. AI-driven population health management	☐	☐	☐	☐	☐	☐
7. Predictive models for resource allocation during emergencies	☐	☐	☐	☐	☐	☐
8. Other, please specify: _____	☐	☐	☐	☐	☐	☐

b. How would you describe your **hospital's** current level of AI implementation in the following operational areas?

	(1) Not Implementing	(2) Exploring	(3) Piloting/Testing	(4) Expanding	(5) Fully Integrated	(0) Don't Know
1. Revenue cycle management (e.g., billing, claims processing)	☐	☐	☐	☐	☐	☐
2. Supply chain optimization	☐	☐	☐	☐	☐	☐
3. Staff scheduling and workforce management	☐	☐	☐	☐	☐	☐
4. Patient flow and demand forecasting	☐	☐	☐	☐	☐	☐
5. Optimizing operational efficiency	☐	☐	☐	☐	☐	☐
6. Other, please specify: _____	☐	☐	☐	☐	☐	☐

Notes: The figure shows the artificial intelligence questions in the FY2021-FY2024 versions of the AHA Annual Survey.

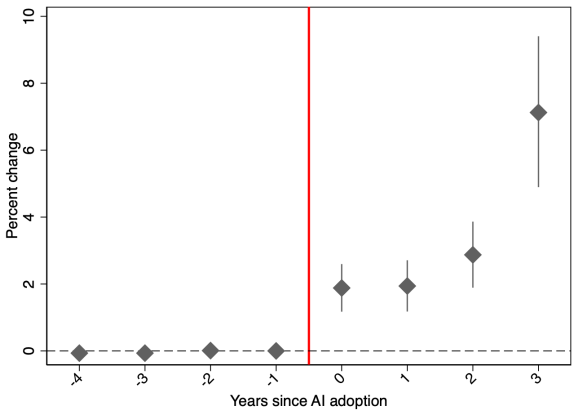
Table A1: Baseline characteristics of Hospital AI adopters and non-adopters, 2017

	AI adopters	**Non-adopters**	**Std. diff.**
N	1,511	2,582	
Finances and Volume			
Net patient revenue (millions $)	339 (483)	139 (215)	0.54
Total expenses (millions $)	342 (491)	139 (227)	0.53
Expenses for clinical activities (millions $)	178 (254)	72 (116)	0.54
Expenses for administrative activities (millions $)	60 (82)	25 (39)	0.55
Discharge equivalents	18,768 (20,242)	9,064 (11,865)	0.58
Labor			
Employees on payroll (FTEs)	1,595 (2,221)	704 (1,084)	0.51
Total paid hours (thousands)	4,094 (4,956)	2,271 (2,681)	0.46
Paid hours, administrative and general (thousands)	416 (641)	244 (315)	0.34
Quality Outcomes			
30-day mortality: heart attack (%)	13.4 (1.2)	13.7 (1.2)	-0.29
30-day mortality: heart failure (%)	11.8 (1.6)	12.1 (1.6)	-0.21
30-day mortality: pneumonia (%)	15.7 (1.9)	16.2 (2.0)	-0.24
30-day mortality: COPD (%)	8.1 (1.1)	8.1 (1.1)	-0.07
30-day mortality: stroke (%)	14.5 (1.7)	14.7 (1.6)	-0.11
30-day mortality: CABG surgery (%)	3.1 (0.8)	3.4 (0.9)	-0.33
30-day all-cause readmission rate (%)	15.3 (0.8)	15.3 (0.8)	-0.11
Patient Experience			
Nurses "Always" communicated well (%)	80.9 (4.3)	82.1 (5.7)	-0.24
Doctors "Always" communicated well (%)	66.6 (7.6)	68.5 (9.3)	-0.22
"Always" received help as soon as wanted (%)	79.8 (4.4)	80.1 (5.5)	-0.06
Hospital rating of 9 or 10 (%)	73.0 (7.1)	72.1 (8.6)	0.12
Would definitely recommend the hospital (%)	73.0 (8.0)	70.9 (9.3)	0.24
Hospital Characteristics			
Beds	221 (241)	114 (139)	0.54
Part of a system	0.87 (0.33)	0.61 (0.49)	0.64
Rural-urban continuum code	2.7 (2.2)	4.1 (2.7)	-0.58
Critical Access Hospital (CAH)	0.19 (0.39)	0.39 (0.49)	-0.46
For-profit	0.07 (0.25)	0.22 (0.41)	-0.43
Major teaching hospital	0.13 (0.34)	0.05 (0.21)	0.31
Medicare casemix index (CMI)	1.68 (0.30)	1.56 (0.32)	0.41
Medicare inpatient discharge share	0.36 (0.15)	0.44 (0.17)	-0.49
Medicaid inpatient discharge share	0.09 (0.09)	0.10 (0.10)	-0.08

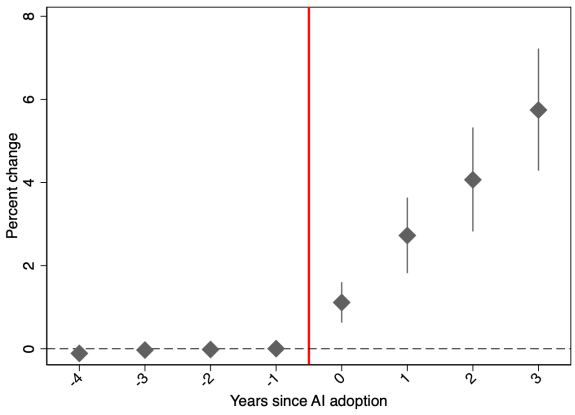
Sources: AHA's Annual Survey (AI adoption), CMS's Hospital Compare (quality measures), Hospital Cost Report data as collected by RAND (all other measures).

Figure A2: Additional finances and volume event studies

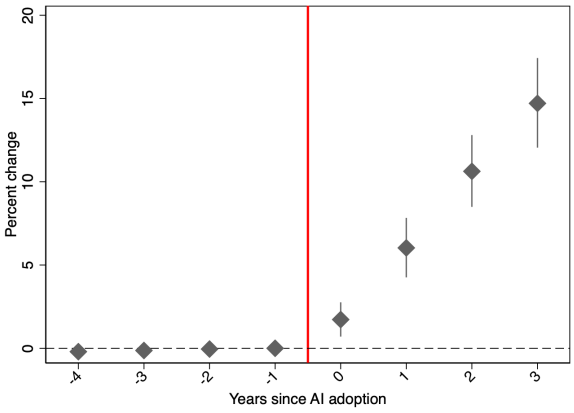
(a) Total expenses



(b) Expenses for clinical activities



(c) Discharge equivalents

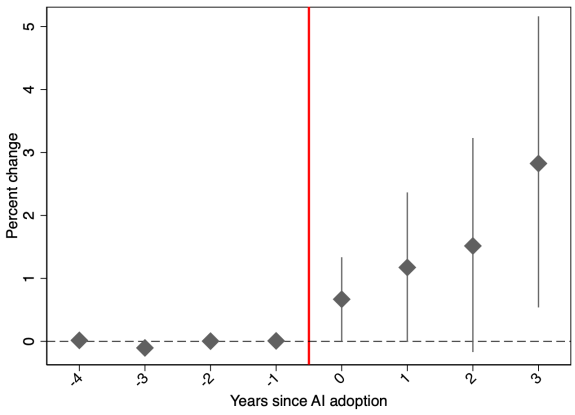


Sources: AHA's Annual Survey (AI adoption) and Hospital Cost Report data as collected by RAND (expenses and discharge equivalents).

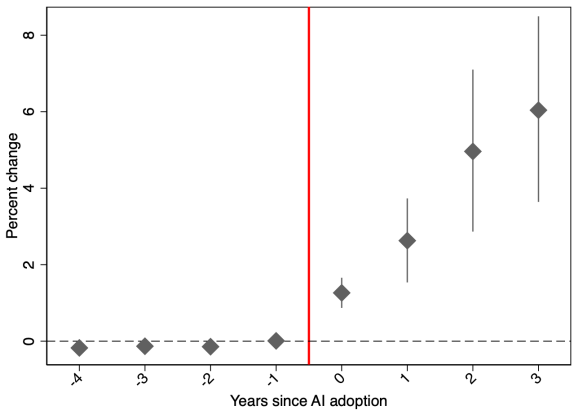
Notes: The figures show estimates from synthetic difference-in-differences event study where the log of each panel's title is the dependent variable. Coefficient estimates have been transformed using the formula exp(coef - 1)*100 to create the percentage effects shown in the figures. The regression model included hospital and year fixed and standard errors were clustered at the hospital level. The error bars represent 95% confidence intervals.

Figure A3: Labor event studies

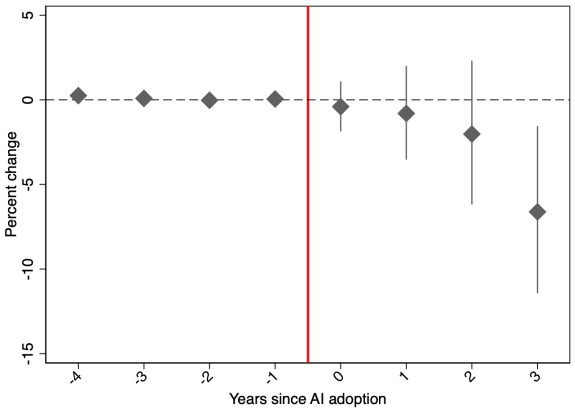
(a) Employees on payroll (FTEs)



(b) Total paid hours



(c) Paid hours, administrative and general

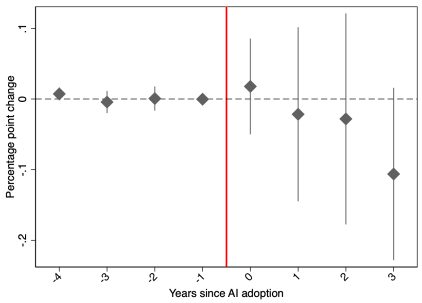


Sources: AHA's Annual Survey (AI adoption) and Hospital Cost Report data as collected by RAND (expenses and discharge equivalents).

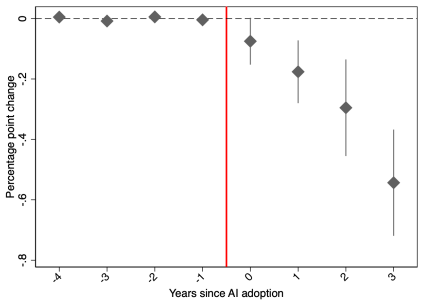
Notes: The figures show estimates from synthetic difference-in-differences event study where the log of each panel's title is the dependent variable. Coefficient estimates have been transformed using the formula exp(coef - 1)*100 to create the percentage effects shown in the figures. The regression model included hospital and year fixed and standard errors were clustered at the hospital level. The error bars represent 95% confidence intervals.

Figure A4: Quality outcomes event studies

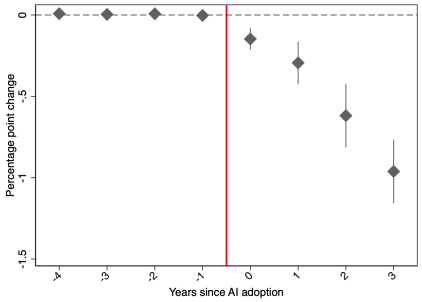
(a) Mortality rate - heart attack



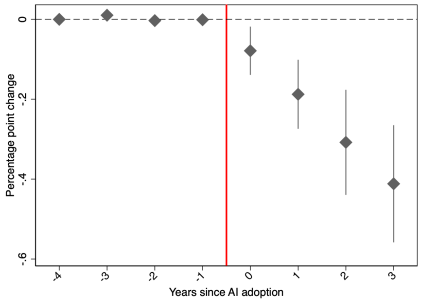
(b) Mortality rate - heart failure



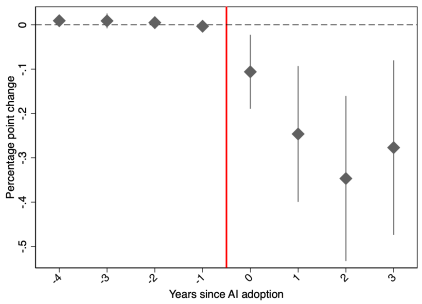
(c) Mortality rate - pneumonia



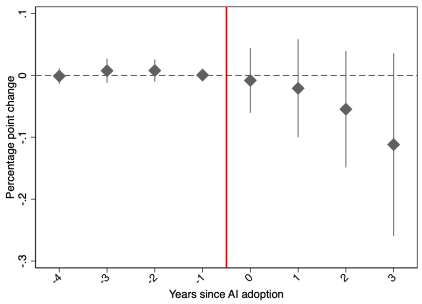
(d) Mortality rate - COPD



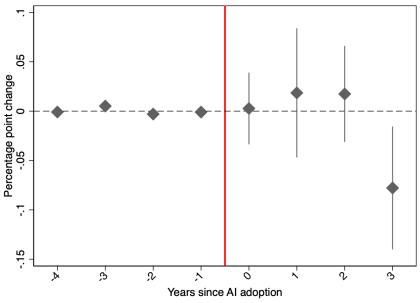
(e) Mortality rate - stroke



(f) Mortality rate - CAPG surgery



(g) All-cause readmission rate

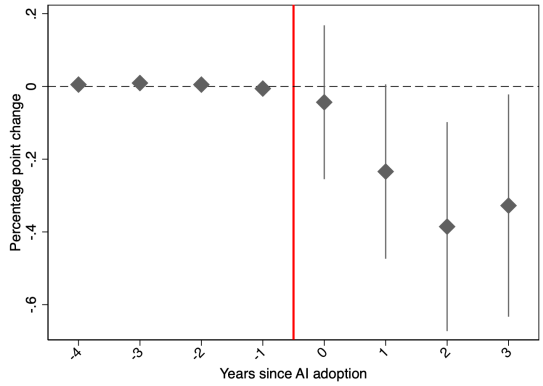


Sources: AHA's Annual Survey (AI adoption) and Hospital Cost Report data as collected by RAND (expenses and discharge equivalents).

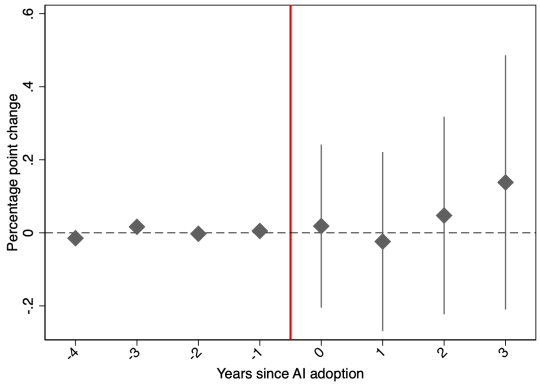
Notes: The figures show estimates from synthetic difference-in-differences event study where each panel's title is the dependent variable. The regression model included hospital and year fixed and standard errors were clustered at the hospital level. The error bars represent 95% confidence intervals.

Figure A5: Patient experience event studies

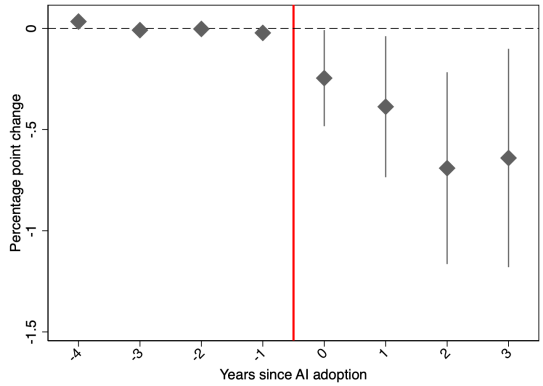
(a) Nurses communicated well



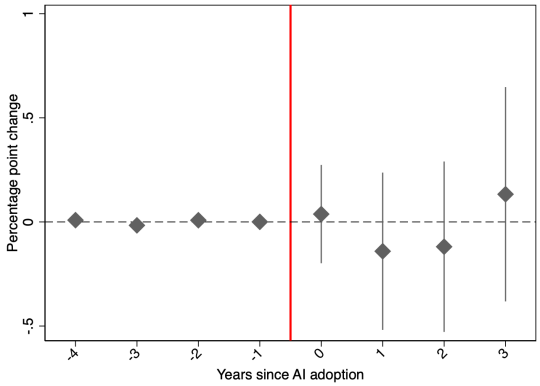
(b) Doctors communicated well



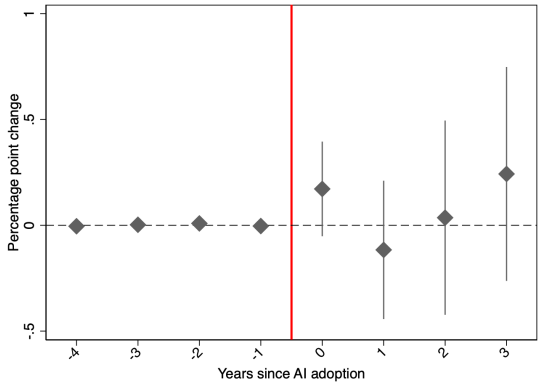
(c) Received help



(d) Hospital rating 9 or 10



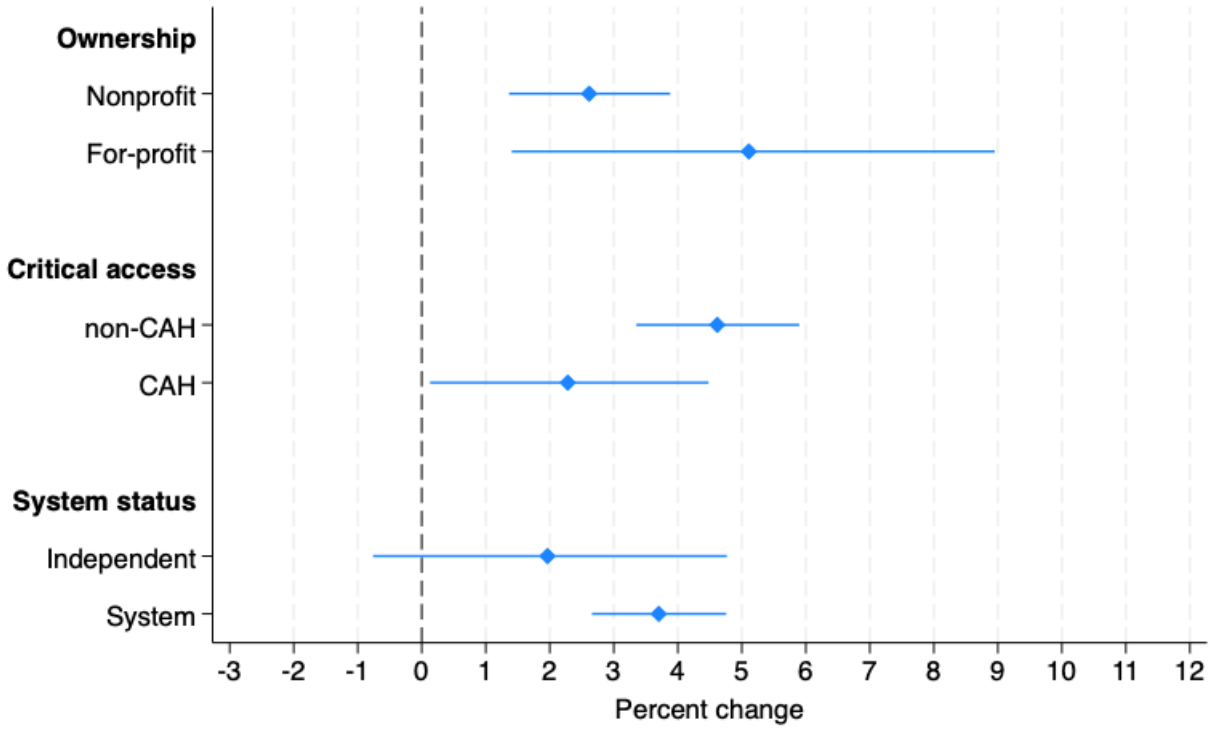
(e) Recommend hospital



Sources: AHA's Annual Survey (AI adoption) and Hospital Cost Report data as collected by RAND (expenses and discharge equivalents).

Notes: The figures show estimates from synthetic difference-in-differences event study where each panel's title is the dependent variable. The regression model included hospital and year fixed and standard errors were clustered at the hospital level. The error bars represent 95% confidence intervals.

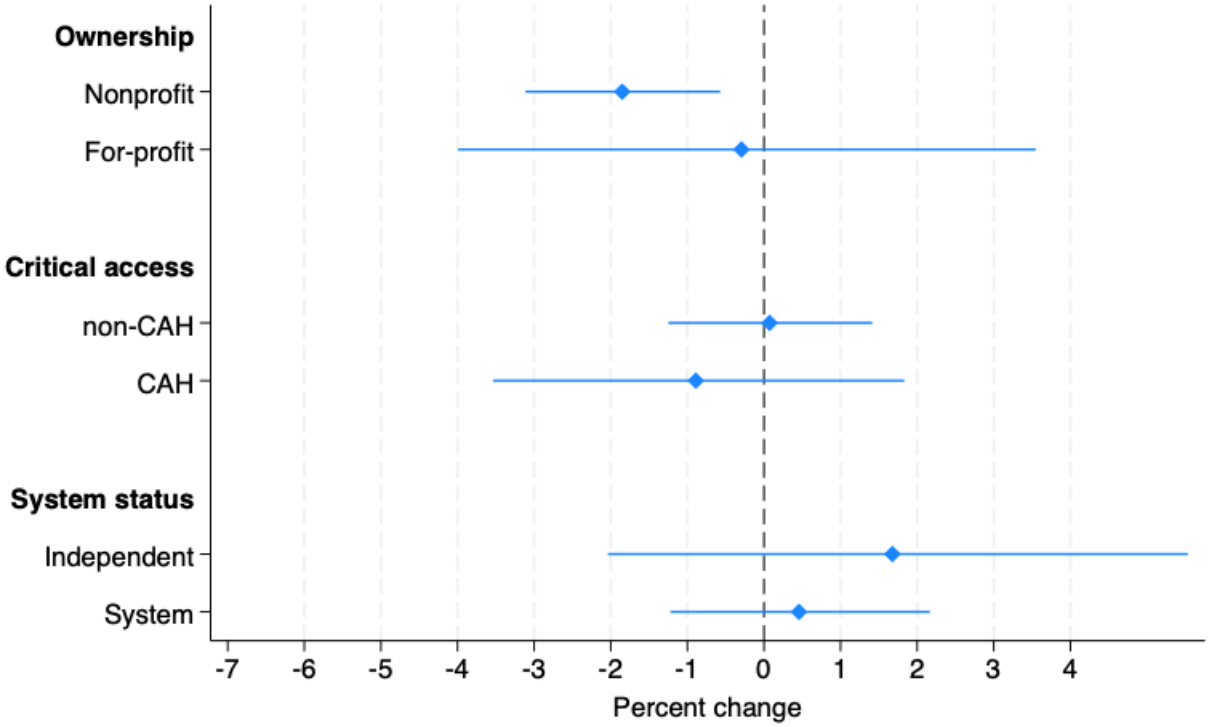
Figure A6: Net patient revenue heterogeneity



Sources: AHA's Annual Survey (AI adoption) and Hospital Cost Report data as collected by RAND (expenses and discharge equivalents).

Notes: The figure shows how the ATT for net patient revenue varies across different subgroups.

Figure A7: Administrative expense heterogeneity

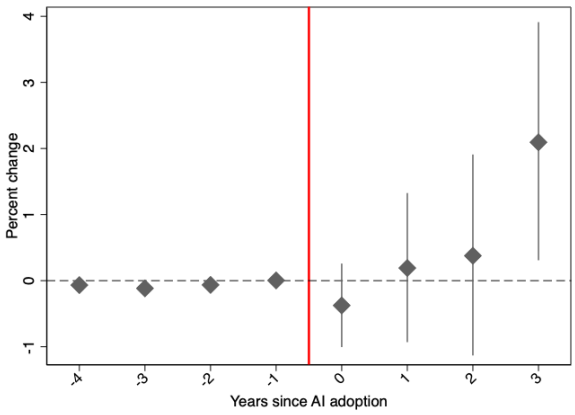


Sources: AHA's Annual Survey (AI adoption) and Hospital Cost Report data as collected by RAND (expenses and discharge equivalents).

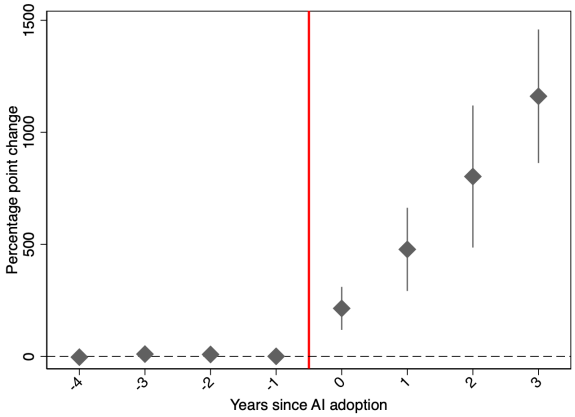
Notes: The figure shows how the ATT for administrative expenses varies across different subgroups.

Figure A8: Additional volume analyses – beds and service type

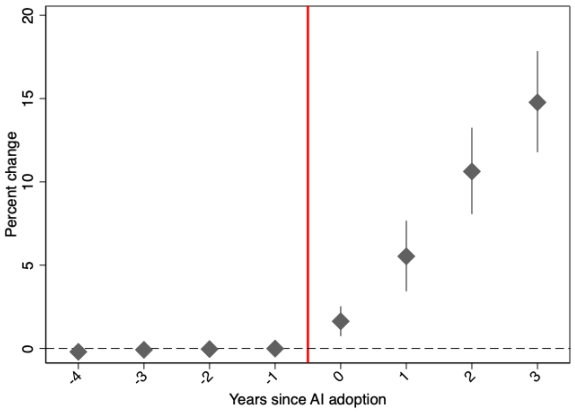
(a) Beds



(b) Inpatient discharges



(c) Outpatient discharge equivalents

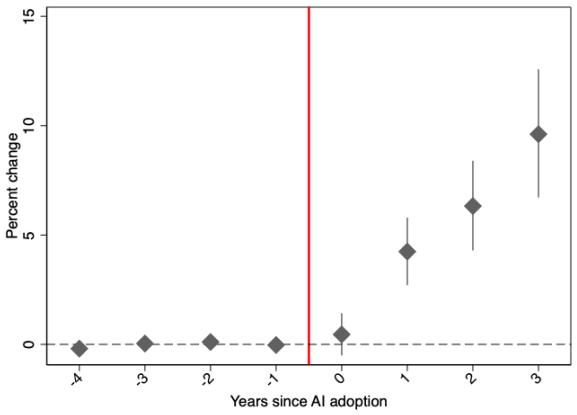


Sources: AHA's Annual Survey (AI adoption) and Hospital Cost Report data as collected by RAND (expenses and discharge equivalents).

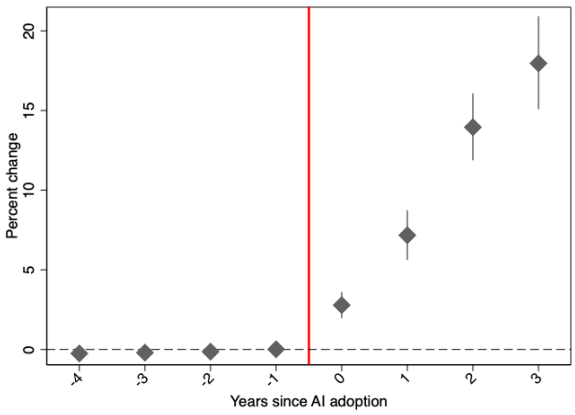
Notes: The figures show estimates from synthetic difference-in-differences event study where the log of each panel's title is the dependent variable. Coefficient estimates have been transformed using the formula exp(coef - 1)*100 to create the percentage effects shown in the figures. The regression model included hospital and year fixed and standard errors were clustered at the hospital level. The error bars represent 95% confidence intervals.

Figure A9: Additional volume analyses – payer

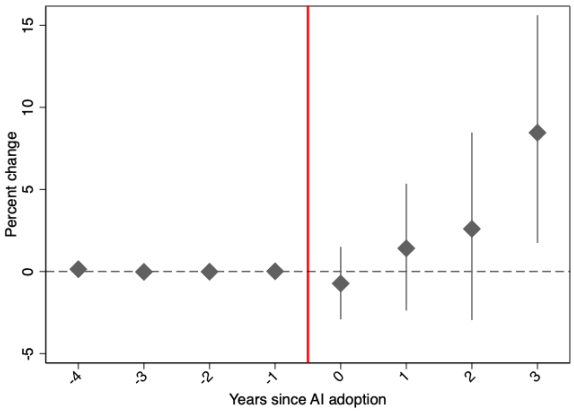
(a) Commercial discharges



(b) Medicare discharges



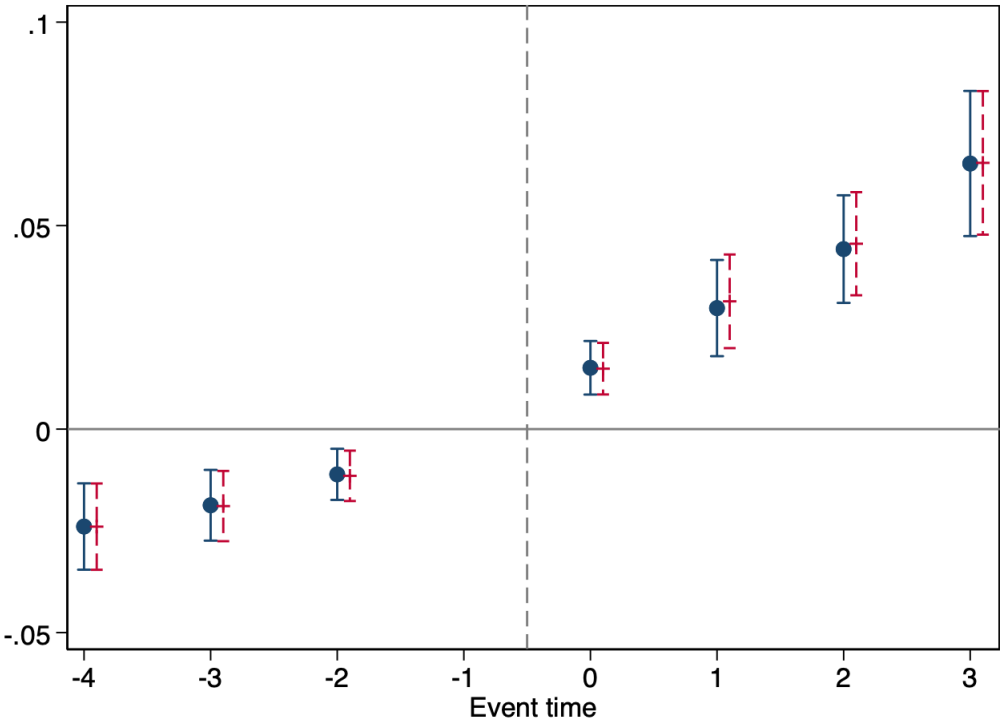
(c) Medicaid discharges



Sources: AHA's Annual Survey (AI adoption) and Hospital Cost Report data as collected by RAND (expenses and discharge equivalents).

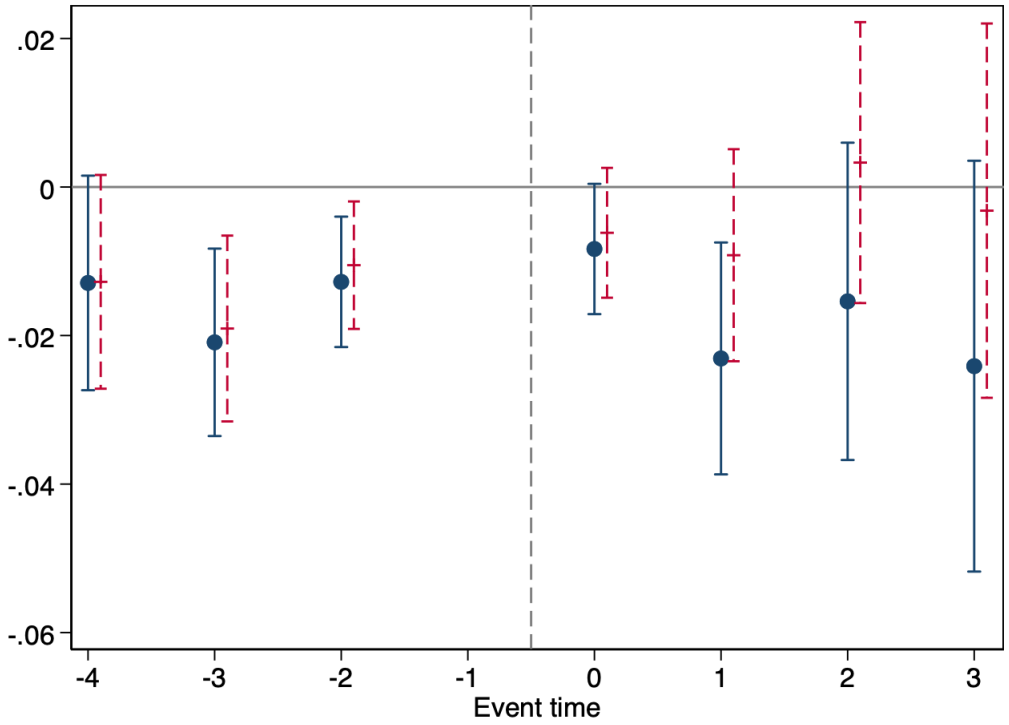
Notes: The figures show estimates from synthetic difference-in-differences event study where the log of each panel's title is the dependent variable. Coefficient estimates have been transformed using the formula exp(coef - 1)*100 to create the percentage effects shown in the figures. The regression model included hospital and year fixed and standard errors were clustered at the hospital level. The error bars represent 95% confidence intervals.

Figure A10: TWFE and Sun-Abraham versions of the change in net patient revenue after AI adoption



Notes: The figure shows the two-way fixed effects (TWFE) (dotted) and Sun-Abraham (solid) versions of SDID event study shown in Figure 2. The dependent variable is ln(net patient revenue).

Figure A11: TWFE and Sun-Abraham versions of the change in administrative expenses after AI adoption



Notes: The figure shows the two-way fixed effects (TWFE) (dotted) and Sun-Abraham (solid) versions of SDID event study shown in Figure 2. The dependent variable is ln(administrative spending).

Details on the aggregate effects calculations

To estimate the aggregate impact of adoption across the U.S. healthcare system, we extrapolated our SDID estimates to the full population of 1,511 adopting hospitals over the 2021–2024 study period. We calculated the average duration of adoption within this four-year window to be 2.5 years.

Finances and Volume Estimates. For financial and volume-based outcomes, we calculated cumulative effects by applying the estimated percentage change (SDID coefficient) to the baseline sample mean (Table S1), scaled by the number of adopters and the duration of adoption.

- *Revenue & Expenses:* Applying the 3.3% increase in net patient revenue to the baseline mean of $339 million yielded an aggregate increase of $42.0 billion. Total expenses followed a similar trajectory, rising by $38.0 billion (based on a 2.9% estimate). Decomposing these costs revealed a shift in resource allocation: while clinical expenses increased by $19.1 billion, administrative expenses decreased by $2.7 billion.
- *Volume:* Operational throughput also expanded; applying the 6.6% estimate to the baseline of 18,768 discharge equivalents suggested a total increase of 4.7 million discharge equivalents across the adopting cohort.

Labor Estimates. To contextualize the increase in labor hours, we converted aggregate hours into employee counts. We estimated a total increase of approximately 490 million labor hours (3.2% effect size $\times$ 4.1 million baseline hours $\times$ 1,511 hospitals $\times$ 2.5 years). Using a standard full-time equivalent (FTE) benchmark of 2,000 hours per year (40 hours/week $\times$ 50 weeks) over the 4-year period, this increase corresponds to approximately 61,000 additional employees.

Mortality Estimates. For mortality outcomes, we employed a population-based market share approach. We first derived annual national hospitalization volumes for Heart Failure, Pneumonia, and COPD by applying historical hospitalization rates (per 100,000) to the average U.S. adult population for 2021–2024. We assumed that adopting hospitals—which account for one-third of total U.S. hospital beds—treated a proportional one-third of these national cases. We then calculated in-hospital death reductions by applying the SDID percentage-point reduction to this specific patient volume.

- *Heart Failure:* With an estimated 4 million national hospitalizations over the study period, the adopting hospitals managed approximately 1.33 million cases. Applying the 0.22 percentage-point reduction in mortality yielded an estimated 2,983 fewer in-hospital deaths.

- *Pneumonia:* Based on 6.4 million national hospitalizations, the adopting hospitals managed approximately 2.13 million cases. The 0.41 percentage-point reduction corresponded to 8,809 fewer in-hospital deaths.

- *COPD:* Based on 2.8 million national hospitalizations, the adopting hospitals managed approximately 0.93 million cases. The 0.21 percentage-point reduction corresponded to 1,953 fewer in-hospital deaths.

- *Stroke:* Based on 2.7 million national hospitalizations, the adopting hospitals managed approximately 0.90 million cases. The 0.21 percentage-point reduction corresponded to 1,987 fewer in-hospital deaths.

Table A2: AI Implementation Status Among General Community Hospitals, 2024

	Not Implementing	Exploring	Piloting/ Testing	Expanding	Fully Integrated	Don't Know	N
Clinical AI Implementation							
AI-Assisted Diagnostics	0.16	0.23	0.16	0.21	0.16	0.08	2,361
Predictive Analytics for Patient Care	0.17	0.22	0.14	0.21	0.18	0.08	2,379
Clinical Decision Support Tools	0.15	0.24	0.10	0.23	0.20	0.08	2,299
AI-Assisted Surgery	0.45	0.23	0.08	0.08	0.03	0.13	2,348
Patient Communication & Education	0.21	0.32	0.20	0.14	0.04	0.09	2,355
Population Health Management	0.26	0.33	0.17	0.06	0.09	0.10	2,380
Resource Allocation During Emergencies	0.45	0.29	0.04	0.03	0.03	0.15	2,348
Other Clinical AI	0.36	0.05	0.06	0.08	0.06	0.38	1,259
Operational AI Implementation							
Revenue Cycle Management	0.19	0.26	0.18	0.23	0.06	0.08	2,356
Supply Chain Optimization	0.28	0.42	0.11	0.05	0.02	0.12	2,343
Staff Scheduling & Workforce Mgmt.	0.26	0.41	0.11	0.09	0.03	0.10	2,287
Patient Flow & Demand Forecasting	0.27	0.36	0.13	0.07	0.06	0.10	2,350
Optimizing Operational Efficiency	0.22	0.31	0.18	0.15	0.05	0.10	2,339
Other Operational AI	0.44	0.05	0.02	0.03	0.03	0.43	1,136

Notes: The table shows how hospitals answered each of 14 AI related questions in the FY2024 Annual Survey. We classified hospitals as adopters in 2024 if they answered "expanding" or "fully integrated" to any of these questions.