CES ifo Working Papers

13086
2026

October 2026

AI and the Boundary of the Firm: Evidence from Radiology

Marco Francesconi, Xiaoyu Xia

CES ifo

Imprint:

CESifo Working Papers

ISSN 2364-1428 (digital)

Publisher and distributor: Munich Society for the Promotion
of Economic Research - CESifo GmbH

Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740

Email office@cesifo.de
https://www.cesifo.org

Editor: Clemens Fuest

An electronic version of the paper may be downloaded free of charge
· from the CESifo website: www.ifo.de/en/cesifo/publications/cesifo-working-papers
· from the SSRN website: www.ssrn.com/index.cfm/en/cesifo/
· from the RePEc website: https://ideas.repec.org/s/ces/ceswps.html

AI and the Boundary of the Firm: Evidence from Radiology*

Marco Francesconi
University of Essex
mfranc@essex.ac.uk

Xiaoyu Xia
University of Essex
xiaoyu.xia@essex.ac.uk

Abstract

Artificial intelligence (AI) may transform labor markets not only by changing workers' tasks but also by reorganizing production. We study radiology, where AI has spread rapidly without increasing interpretations per radiologist. In our model, AI facilitates outsourcing by letting hospitals verify the quality of remotely supplied interpretations. Using US Medicare data for 2013–2024, we document declining local radiology activity in smaller markets despite stable imaging demand. Interpretations shift toward larger markets, and radiologists at documented AI-adopting platforms end up serving more markets. AI can thus reshape firm boundaries and the geographic organization of professional work without replacing workers' tasks.

JEL Classification: J23, L22, L24, O33, I11

Keywords: Artificial intelligence; Labor demand; Technology adoption; Task substitution; Firm boundaries; Outsourcing; Radiology; Health care.

*We are grateful to Daniel Garrett, Piero Gottardi, Christos Makridis, Enrico Moretti, and seminar audiences for insightful comments and discussions.

1 Introduction

The consequences of artificial intelligence (AI) for labor demand are commonly analyzed through the tasks the technology can perform. By automating some tasks and complementing workers in others, AI changes labor demand within and across occupations (Autor, 2015; Acemoglu and Restrepo, 2019a,b). Measures of occupational exposure likewise compare technological capabilities with workers' tasks (Frey and Osborne, 2017; Webb, 2020; Eloundou et al., 2024). Yet AI may also change how production is organized, particularly by improving the monitoring and verification of work supplied across organizational boundaries. Radiology provides a revealing setting for distinguishing these margins. In 2016, Geoffrey Hinton, later a Nobel laureate, famously suggested that we should stop training radiologists because deep learning would soon outperform them at interpreting medical images.[1]

Radiology is among the earliest and most extensively exposed medical specialties. The FDA has authorized more AI-enabled devices for radiology than for all other medical specialties combined (U.S. Food and Drug Administration, 2026), Zhang et al. (2024) count 445 cleared radiology devices between 2008 and 2022, and by 2020 approximately 30 percent of radiologists reported using AI (Allen et al., 2021). Yet AI diffusion has not been accompanied by a contraction in the profession or a discernible increase in interpretations per radiologist. Medicare-billing radiologists continued to grow in number while radiology practices consolidated (Christensen et al., 2024), and large groups expanded their geographic reach (Table 1). Teleradiology also evolved from supplying mainly preliminary overnight interpretations toward providing final interpretations (Steinbrook, 2007; Silva et al., 2013); see Hanna et al. (2020) for an overview of teleradiology's expansion. These developments raise the possibility that AI's most consequential effects in radiology may concern the organization of production rather than the automation of radiologists' tasks.

This paper develops a model in which AI facilitates the contracting of expert services by improving the verification of professional output. Medical image interpretation has features of a credence good: purchasers cannot judge whether the interpretation they received was the one they needed, and may be unable to assess its quality even after the fact (Arrow, 1963; Darby and Karni, 1973; Dulleck and Kerschbamer, 2006). Although digital images can be transmitted at negligible cost, imperfectly observed quality makes contracting for remote interpretations costly. AI-assisted verification can reduce these costs, allowing multi-market providers to pool work across establishments. Because maintaining local interpretation services involves staffing indivisibilities, pooling can reduce the capacity required to serve dispersed demand. This advantage is particularly relevant in thin markets, where declining interpretation volumes may also threaten the viability of local practices. Fixed verification costs also favor adoption by large, multi-client organizations. The model predicts that improved verification can shift production toward multi-market platforms and reduce local radiologist positions without replacing the professional task or necessarily reducing aggregate demand for radiologists.

[1] Hinton made the prediction at the 2016 Machine Learning and the Market for Intelligence conference; see also Langlotz (2019).

We test these predictions using Medicare claims, physician practice directories, and the diffusion of FDA-cleared reader-facing AI between 2013 and 2024. Our empirical strategy interacts national AI diffusion with predetermined differences in markets' radiology interpretation intensity, controlling for market and year fixed effects and differential changes associated with market size. The design identifies changes in the relationship between exposure and local radiology billing as AI becomes more widely available, rather than the causal effect of adoption by individual practices. By 2024, markets one standard deviation more exposed had lost about a fifth more of their locally billing radiologists; the loss is concentrated in markets that began with fewer than ten radiologists and absent in larger markets. Imaging demand measured by patients' residence does not exhibit the same pattern: interpretations billed from smaller markets decline, while larger markets bill more interpretations than their own residents' imaging would generate. Importantly, data on the location where interpretations are performed across Medicare payment localities reveal a similar geographic shift, indicating that the results do not simply reflect changes in billing addresses. The evidence thus points to a reorganization of radiology production across markets rather than a contraction in local imaging demand.

Additional evidence supports the organizational interpretation. Radiology groups expanded their geographic reach, and individual radiologists increasingly served multiple markets without corresponding increases in interpretation throughput, the number of interpretations a radiologist certifies in a year. Documented AI adoption was concentrated among large, multi-market organizations. Within small markets, the exposure-related decline in local radiologist billing is also larger after entry by documented AI-adopting platforms, though not significantly more than after entry by other multi-state groups and not in larger markets. Among radiologists billing from small markets in 2017, approximately 30 percent had ceased Medicare billing by 2024, compared with 21 percent in the largest markets; those who remained also experienced declines in billed interpretation volume. Institutional comparisons rule out the main market-level alternative explanations. Pathologists serve the same hospitals and bill the same payer, so a decline driven by small hospital markets in general—falling rural demand, Medicare Advantage growth, practice consolidation or private-equity acquisition—would appear in their billing too; it does not. Pathology had almost no reader-facing AI and, until 2020, no legal scope for remote certification, and interstate medical licensure compacts, which lowered the cost of practicing across state lines, do not generate similar changes in local radiology billing. Taken together, these findings are consistent with a reduction in the costs of organizing certified work across markets, rather than a large increase in radiologists' task productivity.

A calibrated version of the model assesses the quantitative implications. Combining market-level interpretation volumes, hospital affiliations, and observed geographic reorganization, it reproduces the steep decline in local billing across the smallest markets, including a within-small-market gradient not directly targeted by the calibration. The model implies that approximately a quarter of interpretation production is reorganized across markets, yielding savings equivalent to about 2.5 percent of aggregate reader capacity. By construction, a task-productivity improvement with

the same aggregate saving but fixed organizational boundaries leaves the location of production unchanged. Thus, technologies with similar aggregate efficiency effects can have very different consequences for the location of professional work and the viability of small practices.

Our analysis contributes to the literature on technological change and labor markets by identifying an organizational margin through which AI can affect labor demand. Task-based approaches emphasize substitution and complementarity between technology and workers' activities, while recent empirical studies document productivity gains that are largest for less-experienced workers (Brynjolfsson et al., 2025), heterogeneous employment effects across occupations (Hampole et al., 2025), and task reorganization with no detectable effect on earnings or hours so far (Humlum and Vestergaard, 2025). Evidence from medical applications similarly finds that the effects of AI assistance on diagnostic performance depend on the radiologist and the algorithm (Agarwal et al., 2023; Yu et al., 2024), and that radiologists differ widely in how often they override AI flags (Goldsmith-Pinkham et al., 2026). We complement this literature by showing how AI diffusion can be associated with substantial geographic reallocation of professional work even without a detectable increase in output per professional. The resulting changes in local labor demand need not coincide with the tasks most directly affected by the technology.

Our principal contribution to organizational economics concerns the boundaries of the firm. Since Coase (1937), a central question has been which activities are coordinated within organizations and which are governed through market transactions. The incomplete-contracting literature emphasizes how contractual limitations shape these boundaries (Grossman and Hart, 1986), while research on information technologies shows how monitoring and contractibility affect organizational choices (Baker and Hubbard, 2003, 2004) and how cheaper access to information and cheaper communication shift decision rights within firms (Bloom et al., 2014); Goldfarb and Tucker (2019) survey this work. A related literature models knowledge work as hierarchies in which communication costs determine who solves which problems and where (Garicano, 2000; Garicano and Rossi-Hansberg, 2006; Antràs et al., 2006; Garicano and Rossi-Hansberg, 2015). Verification costs play the analogous role across organizational boundaries here: they govern whether a hospital can contract for interpretations read elsewhere, rather than how expertise is layered within a firm. We develop a mechanism in which AI-assisted verification reduces the cost of contracting for expert services whose quality is difficult to assess. Fixed verification costs favor multi-client providers, connecting the mechanism to economies of scale in geographically expanding service firms (Hsieh and Rossi-Hansberg, 2023); radiology provides a measured case. Our analysis is also related to Ide and Talamàs (2025), who examine AI and the allocation of knowledge-intensive tasks within production hierarchies. Our emphasis instead is on contracting between organizations for work that remains performed and certified by human experts. AI can consequently change firm boundaries and the geographic allocation of production without directly replacing professional tasks.

Finally, we contribute to the literature on the organization of health-care delivery. Existing research examines organizational integration and care utilization (Agha et al., 2023), payment incentives and physicians' employment arrangements (Dranove and Ody, 2019), and hospital ac-

quisitions of physician practices (Capps et al., 2018). We study a different margin: the geographic separation of diagnostic interpretation from the establishments where patients receive care. Unlike research on patients traveling across markets (Dingel et al., 2023), our analysis examines the movement of the professional service itself. This distinction allows us to connect organizational change in health care to the geographic distribution of professional work.

The remainder of the paper is organized as follows. Section 2 presents the motivating facts, and Section 3 develops the model. Section 4 describes the data and empirical strategy. Sections 5 and 6 present the empirical findings, while Section 7 quantifies the implications of organizational reallocation for local practice viability and aggregate reader-capacity requirements and examines counterfactual allocations. Section 8 concludes. The Online Appendix provides supplementary information on the data and exposure measure, additional empirical results, and details of the quantitative model and counterfactual exercises, and lists the data and documentary sources cited in the text.

2 What Did AI Change in Radiology?

Three facts motivate our analysis. First, AI entered radiology without removing the requirement that radiologists certify final interpretations. Second, its diffusion was not accompanied by a discernible increase in interpretations per radiologist. Third, radiology production became increasingly organized across establishments and geographic markets. Together, these facts motivate an examination of AI's organizational effects alongside its effects on task productivity.

2.1 Fact 1: AI Assists, Radiologists Certify

AI entered radiology early and at considerable scale. FDA-cleared reader-facing devices increased from 10 in 2017 to 239 by 2025,[2] and clinical adoption followed, particularly among larger practices (Allen et al., 2021). Yet AI assistance did not eliminate professional responsibility for final interpretations. Diagnostic imaging examinations culminate in reports completed under applicable professional and regulatory requirements (American College of Radiology, 2025) and Medicare's physician-interpretation rule (Code of Federal Regulations, 2026c), professional society guidance assigns final responsibility for AI-assisted interpretations to the radiologist (Geis et al., 2019), and FDA-cleared radiology AI devices function as aids to, not substitutes for, the interpreting clinician (Zhang et al., 2024); none independently replaced the radiologist responsible for the final report during our sample period.

Among the 249 reader-facing devices in our classification, 118 support triage, 99 second reading, and 32 quantification. Triage tools screen problems before they are routed, the organizational role that Carmona and Laohakunakorn (2024) analyze. None of the 249 is explicitly classified as a verification technology. These applications prioritize examinations, flag suspected abnormali-

[2]We classify reader-facing devices using the US Food and Drug Administration's AI-Enabled Medical Device List and associated clearance documents. The list is periodically updated (U.S. Food and Drug Administration, 2026).

ties, automate measurements, or provide additional information to radiologists (Allen et al., 2021; Zhang et al., 2024), and remain narrow and task-specific (Dogra et al., 2025). Although primarily designed to assist clinical work, some functions may also provide information useful for assessing interpretation quality. Whether this reduces the cost of monitoring expert work is an economic question, not something established by FDA classifications.

The continued requirement for professional certification does not preclude productivity gains. AI may reduce interpretation time or effort, allowing more certified reads per radiologist or the same volume with less labor, as in task-based models of automation (Acemoglu and Restrepo, 2019b). Clinical evaluations accordingly examine reading time, turnaround time, and diagnostic performance (Wenderott et al., 2024; Gu et al., 2026). Whether these improvements translate into higher certified output per radiologist is the subject of our second fact.

2.2 Fact 2: Reads per Radiologist Did Not Rise

The aggregate data show no discernible increase in interpretation throughput, interpretations certified per radiologist, as reader-facing AI diffused. Between 2013 and 2024, the number of radiologists billing Medicare rose by 16 percent, while billed imaging interpretations declined by 3 percent. Within one health system, Goldsmith-Pinkham et al. (2026) likewise find order-to-diagnosis time roughly unchanged after an AI rollout, and their reported counts imply volume concentrated among fewer signing radiologists.[3] The decline in interpretations reflects the shrinking Medicare fee-for-service population, which fell by 10 percent as Medicare Advantage expanded: interpretations per fee-for-service beneficiary rose by about 9 percent, as did imaging events per beneficiary in the Geographic Variation data, which include studies not billed by radiologists. Interpretations per billing radiologist consequently fell by 16 percent: flat through the rapid expansion of FDA clearances after 2017, they dropped in the pandemic and have not recovered, while the number of billing radiologists kept growing (Figure 1).

The absence of an upward break contrasts with the labor-saving potential of task automation (Acemoglu and Restrepo, 2019a,b). Task-level forecasts also anticipate only modest aggregate productivity gains from AI within a decade (Acemoglu, 2025). Although individual studies often report reading-time savings, pooled estimates are insignificant and heterogeneity is high (Wenderott et al., 2024), and effects on workflow and diagnostic performance are similarly application-specific (Gu et al., 2026). Moreover, the national series are descriptive: changes in imaging demand, case mix, working hours, or workforce composition could obscure productivity improvements for particular readers or tasks. We therefore examine within-reader and within-firm changes around dated AI adoption in Subsection 5.5, including comparisons between AI-covered and other tasks. Those analyses likewise provide little evidence of increased certified output per radiologist. The

[3] Reading-time gains in reader studies, about a quarter in a meta-analysis of mostly retrospective designs (Chen et al., 2024), need not raise interpretations per radiologist. A radiologist's annual interpretations are set by the volume routed to the reader, which did not rise (Figure 1); deployed tools have not been found to change time per scan (Schmuelling et al., 2021); concentrating reads among fewer readers raises their counts but not the profession's; and the capacity that pooling releases appears as positions not needed rather than as more reads per reader (Section 7).

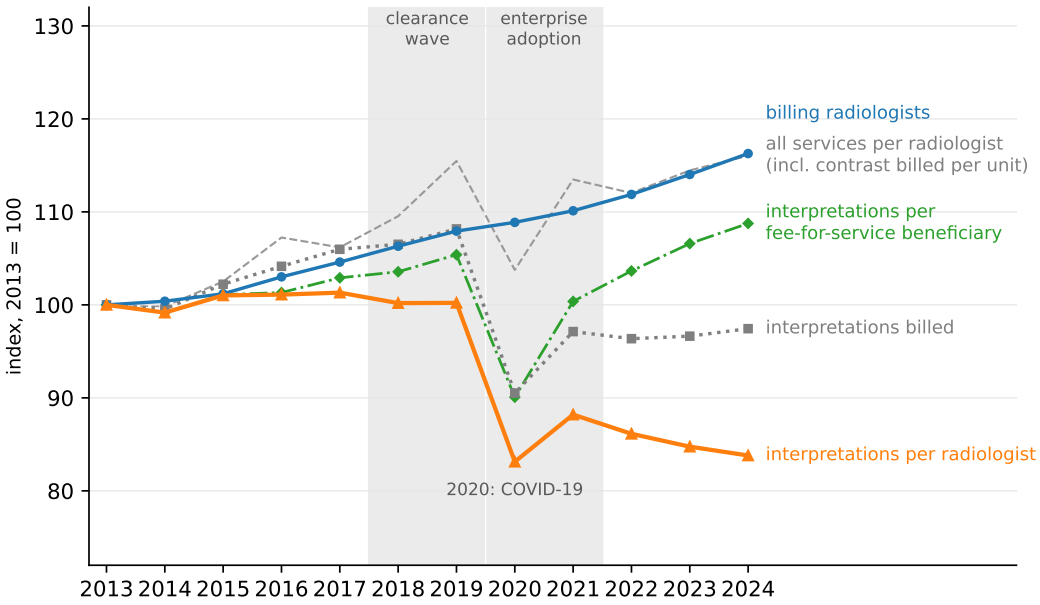


Figure 1: National radiology activity and interpretation throughput, 2013–2024

NOTE.—Series show Medicare-billing radiologists, imaging interpretations billed by radiologists, interpretations per billing radiologist, and interpretations per fee-for-service beneficiary, indexed to 2013=100. Interpretations include radiology procedure codes 70010–79999 and pre-2018 mammography G-codes across all settings, excluding the computer-aided-detection add-on codes bundled into the mammography codes from 2017. The thin dashed series, all Medicare services per radiologist, includes contrast agents. Shaded bands indicate the 2018–19 clearance wave and 2020–21 enterprise deployments. The 2020 decline reflects the pandemic.
SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files and Medicare Monthly Enrollment (July).

evidence does not rule out task-specific productivity gains, but it reveals no detectable increase in aggregate interpretation throughput during AI diffusion.

2.3 Fact 3: Production Reorganizes

Radiology production became increasingly organized across geographic markets. Between 2017 and 2025, the share of radiologists affiliated with groups operating in at least five states rose from 9 to 47 percent, while the number of such groups increased from 79 to 448 (Table 1). The share of radiologists listing practice locations in multiple local markets also increased, from 20 to 36 percent. This geographic expansion accompanied broader consolidation: between 2014 and 2023, the number of Medicare-enrolled radiologists rose by 17 percent while the number of practices fell by 15 percent. Practices with at least 100 radiologists increased 4.5-fold, whereas those with three to nine radiologists declined by one-third (Christensen et al., 2024). Private-equity acquisitions also continued during this period (Khunte and Singh, 2025). Consolidation was thus accompanied by a substantial expansion in the geographic reach of radiology organizations.

Falling image-transmission costs alone cannot readily explain this reorganization. Digital imaging had made remote interpretation feasible well before the recent AI wave: teleradiology was already widely used by 2007 (Steinbrook, 2007), and professional guidance addressed both prelim-

inary and final remote interpretations (Silva et al., 2013). Historically, however, remote services often supplied preliminary or overnight reads, while final interpretation remained with the local practice (Steinbrook, 2007); professional standards continue to distinguish preliminary from definitive reports (American College of Radiology, 2025). The ability to transmit images therefore did not eliminate the organizational advantages of local provision. If remote pooling was technically feasible, why did final interpretations remain substantially organized through local practices (in 2017, 9 percent of radiologists belonged to groups spanning five or more states; Subsection 4.3), and what subsequently made multi-market provision more attractive? These questions point beyond transmission costs toward the costs of contracting for certified interpretations across organizational boundaries.

2.4 The Missing Margin

The three facts point to an organizational margin of adjustment that neither task substitution (Acemoglu and Restrepo, 2019b) nor declining communication costs (Steinbrook, 2007) alone readily explains. AI diffusion coincided with expanding multi-market provision despite little detectable increase in interpretations per radiologist and the long-standing feasibility of remote image transmission. These patterns are descriptive, however, and consolidation predates the recent AI wave. Hospital ownership and private-equity acquisitions may also contribute to changes in organizational structure (Rosenkrantz et al., 2020; Khunte and Singh, 2025). An AI-related explanation must therefore identify a contracting friction that the technology could plausibly relax, rather than attributing the broader consolidation trend to AI.

One such friction is the difficulty of assessing expert work supplied outside the organization. Even when remote providers can pool professional capacity across clients, imperfectly observed quality may make outsourcing costly relative to local provision. Improvements in monitoring and contractibility can alter this make-or-buy decision without changing the production technology (Baker and Hubbard, 2003, 2004). Although the AI devices in our exposure measure are not classified as verification technologies, some may generate information useful for evaluating remotely supplied interpretations. If that information reduces monitoring costs, multi-client providers can exploit pooling opportunities and spread fixed verification costs across establishments. Section 3 formalizes this mechanism and derives its implications for organizational boundaries and the geographic allocation of professional work.

3 AI, Verification, and the Boundary of the Firm

We develop a model in which AI affects the organization of expert services through task productivity and output verification. Indivisible professional capacity creates gains from pooling across establishments, but imperfect observability of expert effort makes remote contracting costly. AI-assisted verification can relax this friction and facilitate pooling without increasing individual productivity. The framework builds on theories of firm boundaries and information-dependent contracting

(Coase, 1937; Grossman and Hart, 1986; Baker and Hubbard, 2003, 2004), relational and subjective-performance evaluation (Baker et al., 2002; Levin, 2003; MacLeod, 2003), and fixed-cost technologies and firm scale (Hsieh and Rossi-Hansberg, 2023; De Ridder, 2024).

3.1 Local Production and the Gains from Pooling

Establishments and demand. A continuum of establishments, indexed by m, generates flows of certified expert tasks ("reads") v_m, distributed according to the measure G over $v \in [\underline{v}, \bar{v}]$. Aggregate demand is $V \equiv \int v\, dG(v)$ and grows exogenously at rate g. In radiology, reads are generated by upstream clinical demand, which we take as given. Each read is reimbursed at an administered price p, the Medicare physician fee schedule (Centers for Medicare & Medicaid Services, 2026h), and requires a certified interpretation. We abstract from any response of imaging demand to lower interpretation costs.

Readers. Certified interpretations are supplied by licensed experts ("readers"), radiologists in our application. Each reader can complete at most $\bar{h}$ reads per period, has reservation utility $\bar{U}$, and incurs private cost $C > 0$ from diligent effort; shirking entails no effort cost and produces deficient reads. The competitive cost of a fully utilized diligent reader is $W \equiv \bar{U} + C$, implying labor cost per read $c_L \equiv W/\bar{h}$. Every final interpretation requires certification by a licensed reader, although AI may augment reader productivity as considered in Subsection 3.3.

In-house production and the indivisibility penalty. In-house staffing is indivisible: readers are employed in integer numbers, and establishments may require a minimum complement $n_{\min} \geqslant 1$ for clinical coverage. Thin local demand constrains the organization of professional services more generally: in law, the extent of the market bounds specialization and hierarchy (Garicano and Hubbard, 2007). The number of readers required for v interpretations is

$$n(v) = \max\left\{n_{\min}, \left\lceil \frac{v}{\bar{h}} \right\rceil\right\}. \tag{1}$$

The average cost of in-house production is

$$C_I(v) = \frac{W n(v)}{v} = c_L[1 + \pi(v)], \qquad \pi(v) \equiv \frac{n(v)\bar{h}}{v} - 1 \geqslant 0. \tag{2}$$

The *indivisibility penalty* $\pi(v)$ measures the excess staffing capacity required by local production relative to the capacity needed to complete the same volume of interpretations under perfectly divisible staffing. This excess may reflect minimum coverage requirements, scheduling constraints, or other indivisibilities, rather than radiologists being literally idle. For $v \leqslant n_{\min}\bar{h}$, the penalty equals $n_{\min}\bar{h}/v - 1$; at higher volumes, integer staffing generates a sawtooth pattern with a declining envelope. Pooling reads across establishments can therefore reduce required staffing capacity without changing the effort needed for individual interpretations. Using the annual reader capacity

assumed in Section 7, a hospital performing 3,000 reads annually requires 33 percent more capacity than its volume, compared with 10 percent for one performing 40,000 reads.

3.2 The Contracting Friction

The quality of expert work may be harder to assess when interpretations are supplied remotely. Local production facilitates monitoring through direct interaction, peer review, and clinical feedback; professional guidance on teleradiology accordingly stresses communication with clinicians and formal peer review with longitudinal follow-up (Silva et al., 2013), consistent with deficiencies in remotely supplied reports becoming apparent mainly through subsequent review.

Assumption 1 (Local observability). A deficient in-house read is detected with probability δ_I, while a deficient remote read is detected with probability $\delta_R < \delta_I$ before effective contractual recourse expires. We normalize $\delta_I = 1$ and write $\delta \equiv \delta_R \in (0, 1)$.

The assumption concerns actionable monitoring, not perfect observation of local medical care. A platform pools interpretations across establishments, eliminating local staffing indivisibilities, but remote provision incurs a per-read coordination cost $t_0 \geqslant 0$, including worklist integration and credentialing. We assume limited liability for readers and platforms, and that the platform cannot enforce a more informative performance signal against its readers than the buyer can observe. The strongest available sanction is therefore withholding payment when a deficiency is detected.

Consider a remote reader paid W_R if no deficiency is detected and zero otherwise. Diligence yields $W_R - C$, while shirking yields $(1 - \delta)W_R$. Incentive compatibility and participation require $\delta W_R \geqslant C$ and $W_R - C \geqslant \bar{U}$, respectively.

Lemma 1 (Agency rent). The minimum compensation required to induce diligence under remote production is

$$W_R^*(\delta) = \max\left\{W, \frac{C}{\delta}\right\}, \tag{3}$$

with agency rent $R(\delta) = \max\{0, C/\delta - W\}$. Without productivity augmentation, the agency rent per read is $\rho(\delta) \equiv R(\delta)/\bar{h}$, and the unit cost of remote production is

$$C_R(\delta) = c_L + t_0 + \rho(\delta). \tag{4}$$

The agency rent decreases strictly with δ whenever positive and vanishes for $\delta \geqslant \widehat{\delta} \equiv C/W$.

Limited liability and imperfect performance signals can generate agency rents (Innes, 1990), while relational and subjective-performance contracts provide related mechanisms (Baker et al., 2002; Levin, 2003; MacLeod, 2003). An establishment compares $C_I(v) = c_L[1+\pi(v)]$ with $C_R(\delta) = c_L + t_0 + \rho(\delta)$: pooling favors remote provision, while coordination costs and agency rents favor local production. Reputation need not eliminate the latter friction when deficiencies are difficult to observe or credibly communicate (Klein and Leffler, 1981; Levin, 2003; MacLeod, 2003).

3.3 AI: Productivity versus Verification

AI capability κ can affect expert services through two distinct channels.

Task productivity. Let $a(\kappa) \geqslant 1$, with $a'(\kappa) \geqslant 0$, denote the increase in reader capacity generated by AI. A reader can then complete $a(\kappa)\bar{h}$ interpretations per period, reducing labor cost per read to $c_L/a(\kappa)$. An establishment producing v reads in-house requires

$$n(v,\kappa) = \max\left\{n_{\min}, \left\lceil \frac{v}{a(\kappa)\bar{h}} \right\rceil\right\},$$

at unit cost $C_I(v,\kappa) = Wn(v,\kappa)/v$. The baseline model corresponds to $a(\kappa) = 1$.

Output verification. AI may also improve the detection of deficient remote interpretations. Let $\delta(\kappa, v)$ denote the probability of detecting a deficiency in time for contractual recourse, where $\delta(\kappa, v) \geqslant \delta_0(v)$ and $\partial\delta(\kappa, v)/\partial\kappa > 0$. The baseline probability $\delta_0(v)$ depends on client volume, reflecting the information available from repeated interactions.

Assumption 2 (Statistical detection). Without verification technology, $\delta_0'(v) > 0$: deficiencies are more readily detected in larger client relationships, which generate more information for formal or relational enforcement.

This assumption captures the informational disadvantage of thin client relationships (e.g., Levin, 2003; MacLeod, 2003).

Assumption 3 (Verification technology). An investment $\kappa > 0$ has three properties:

1. It raises detection and provides a common monitoring standard across clients: $\delta(\kappa, v) = \delta(\kappa) \geqslant \delta_0(v)$.
2. It requires a firm-level fixed cost $F(\kappa) > 0$, while monitoring capacity is non-rival across the firm's reads and client relationships.
3. It generates signals that buyers can observe and use for contractual enforcement. Information available only to the supplier does not improve the buyer's effective detection probability.

These properties allow multi-client platforms to spread verification costs across establishments while reducing the informational disadvantage of remote contracting. Deployed tools generate such signals: Goldsmith-Pinkham et al. (2026) observe, for most scans in one health system after deployment, whether the AI flag and the radiologist's diagnosis agreed and, for the subset with follow-up imaging, how disagreements resolved. The mechanism resembles the monitoring role of information technology in Baker and Hubbard (2003, 2004), but concerns expert output rather than worker inputs. Unlike technologies that lower the marginal cost of replicating services across establishments (Hsieh and Rossi-Hansberg, 2023), verification facilitates expansion by reducing contracting costs without necessarily changing the professional task. The model allows both channels; the following analysis isolates their organizational implications.

3.4 Verification and the Make-or-Buy Boundary

To isolate the organizational channel, hold reader productivity fixed and consider an improvement in verification. By Lemma 1, the agency rent $\rho[\delta(\kappa)]$ decreases with κ whenever the incentive

constraint binds. Under Assumption 3, post-adoption detection is common across clients; write $\rho(\kappa) \equiv \rho[\delta(\kappa)]$. An establishment purchases remote interpretations whenever $C_R(\kappa) \leqslant C_I(v)$, or equivalently,

$$\pi(v) \geqslant \pi^*(\kappa) \equiv \frac{t_0 + \rho(\kappa)}{c_L}. \tag{5}$$

Improved verification thus lowers the minimum indivisibility penalty required for outsourcing.

Proposition 1 (Verification and the boundary of the firm) *With reader productivity fixed, an establishment outsources if and only if $\pi(v) \geqslant \pi^*(\kappa)$. Improved verification weakly expands the outsourcing set $S(\kappa) = \{v : \pi(v) \geqslant \pi^*(\kappa)\}$, strictly so for establishments whose indivisibility penalties lie between the old and new thresholds. Once $\delta(\kappa) \geqslant \widehat{\delta}$, the agency rent vanishes and the outsourcing set reaches $\{v : \pi(v) \geqslant t_0/c_L\}$.*

The result does not require $\pi(v)$ to decline monotonically with volume: integer staffing generates local discontinuities, but establishments with sufficiently large indivisibility penalties are the first to find outsourcing attractive. Verification does not create the gains from pooling; it reduces the informational cost of realizing them.

The mechanism also helps explain the historical reliance on preliminary teleradiology. An establishment could purchase a remote interpretation while retaining local certification, effectively using a second professional assessment in place of remote verification. If both interpretations require a full unit of reader time, the approximate unit cost is $C_R^{dup} = 2c_L + t_0$. Such duplication is more attractive when maintaining local coverage is especially costly, as during overnight shifts (Steinbrook, 2007). Improved verification can instead facilitate remote provision of final interpretations without a duplicative local read.

Equation (5) characterizes the organizational response conditional on verification adoption. In Section 7, we implement the organizational margin through a reduced-form relationship between staffing indivisibilities and cross-market affiliations. The calibration maps observed organizational changes into production reallocation without separately identifying detection probabilities, agency rents, or fixed adoption costs.

3.5 Labor Reallocation without Task Substitution

When an establishment switches from in-house to pooled interpretation, its reads remain certified by licensed readers, but some of the staffing capacity previously maintained to provide interpretations locally can be shared across establishments. This does not imply that all on-site radiology responsibilities disappear: clinical procedures and other services requiring physical presence may continue to require local staffing. The model isolates the interpretation component of professional work, for which an establishment producing v reads requires $n(v)$ readers under local production, whereas pooled production requires $v/\bar{h}$ units of reader capacity. The aggregate capacity saving from outsourcing is then

$$E(\kappa) = \int_{S(\kappa)} \left[n(v) - \frac{v}{\bar{h}} \right] dG(v) = \frac{1}{\bar{h}} \int_{S(\kappa)} \pi(v) v \, dG(v) \geqslant 0. \tag{6}$$

The saving is strictly positive whenever the outsourcing set contains a positive measure of establishments with unused capacity.

Proposition 2 (Labor reallocation without task substitution) *Holding reader productivity fixed, an improvement in verification that expands the outsourcing set:*

1. *preserves the requirement that every interpretation be certified by a licensed reader;*
2. *transfers outsourced interpretations from in-house production to pooled suppliers;*
3. *weakly reduces aggregate reader capacity required for a given volume of interpretations, with savings given by Equation (6);*
4. *generates larger capacity savings where newly outsourced establishments have greater aggregate unused staffing capacity.*

Unlike task substitution, pooling leaves labor requirements per completed interpretation unchanged. It reduces aggregate required capacity by eliminating underutilized local staffing, while shifting interpretation work toward remote suppliers. Consequently, local positions may decline even as aggregate professional employment grows, provided demand growth offsets the capacity savings. Multi-market platforms can also serve more establishments with the same reader capacity, increasing geographic coverage per reader without increasing interpretations per reader.

A second indivisibility: practice exit. Reorganization may render a local practice unviable when its retained interpretation volume falls below minimum staffing requirements. Consider a local production unit initially supplying v reads, of which v' remain after some production moves elsewhere. At reimbursement p per read, continued production requires $pv' \geqslant Wn(v')$, or equivalently $\pi(v') \leqslant p/c_L - 1$. When minimum staffing binds, $n(v') = n_{\min}$, the practice cannot cover its labor costs if

$$v' < n_{\min}\bar{h}\,\frac{c_L}{p}.$$

Exit reallocates the remaining interpretations to pooled suppliers, potentially amplifying the initial decline in local production and positions. This mechanism is particularly relevant in thin markets. Section 7 implements it through a hospital-level viability threshold, assuming a minimum staffing complement of one reader and $c_L/p = 1$, without explicitly modeling individual practices' contractual portfolios.

3.6 Adoption, Scale, and Market Structure

Verification requires a fixed investment $F(\kappa)$. Consider a platform serving clients $\mathcal{M}$, with volumes $\{v_m\}_{m\in\mathcal{M}}$ and total monitored volume $Q = \sum_{m\in\mathcal{M}} v_m$. For client m, adoption reduces the agency rent per read by $\Delta\rho_m(\kappa) \equiv \rho[\delta_0(v_m)] - \rho[\delta(\kappa)] \geqslant 0$. Holding the client portfolio fixed, the net operating-cost saving is

$$\Pi(\kappa;\mathcal{M}) = \sum_{m\in\mathcal{M}} \Delta\rho_m(\kappa)v_m - F(\kappa). \tag{7}$$

Adoption is cost-saving when $\Pi(\kappa; \mathcal{M}) \geqslant 0$. If the per-read saving is approximately common across clients, this reduces to $Q \geqslant Q^*(\kappa) \equiv F(\kappa)/\Delta\rho(\kappa)$.

Proposition 3 (Scale and client composition) *Under Assumptions 2 and 3, verification adoption has two properties:*

1. Scale. *Replicating a portfolio of client relationships increases the gross benefit of verification proportionally while leaving its fixed cost unchanged.*
2. Client fragmentation. *Holding total volume fixed, dividing an existing client relationship into smaller relationships weakly increases the gross benefit of verification, and increases it strictly wherever the baseline agency rent declines with client volume.*

The scale result follows from the non-rival nature of verification. For client fragmentation, splitting a relationship of volume v into smaller relationships increases their baseline agency rents because detection improves with client volume under Assumption 2. Post-adoption agency rents are common across clients, so the gross saving from verification weakly increases. Verification thus favors platforms serving many individually thin relationships, even conditional on total volume. This distinguishes it from a generic fixed-cost technology whose benefits depend only on aggregate output.

As verification expands the outsourcing set, the volume supplied by platforms, $V_S(\kappa) \equiv \int_{S(\kappa)} v\, dG(v)$, increases. Platforms can spread fixed verification costs across more reads and client relationships, potentially favoring larger multi-client suppliers and greater concentration in remote provision (Hsieh and Rossi-Hansberg, 2023; De Ridder, 2024). These scale economies arise from governing production across organizational boundaries rather than necessarily lowering the cost of performing individual interpretations.

Proposition 3 does not say who can pay $F(\kappa)$. A group can reach Q^* either by adding clients one contract at a time or by acquiring groups that bring their contracts with them. The vRad case illustrates the pairing: under Mednax it built its own worklist models while local groups kept the hospital contracts (MEDNAX, Inc., 2015, 2017); after Radiology Partners bought it in 2020 (MEDNAX, Inc., 2020), AI was deployed across the combined firm's hospital contracts and discrepancy reports went to client hospitals (Radiology Partners, 2021; vRad (Virtual Radiologic), 2021). This example illustrates how investments in verification and the expansion of client portfolios can reinforce one another. Acquisitions may allow platforms to spread fixed verification costs across more contracts, while improved verification may increase the value of serving geographically dispersed clients. Private-equity acquisitions account for a substantial share of recent consolidation in radiology (Khunte and Singh, 2025), although the model does not require any particular source of capital or imply that AI adoption caused these acquisitions.

3.7 Distinguishing Verification from Alternative Technologies

Three technological channels can affect the organization of expert services. *Task productivity*, operating through $a(\kappa)$, increases interpretations per reader and reduces labor requirements per

read, but may also alter outsourcing incentives when local staffing is indivisible. *Generic scale technologies*, such as centralized billing, scheduling, or credentialing, can encourage consolidation by lowering the cost of serving additional markets rather than by changing the professional task (Hsieh and Rossi-Hansberg, 2023; De Ridder, 2024). *Verification*, operating through $\delta(\kappa)$ and the agency rent $\rho(\kappa)$, reduces the cost of contracting for remotely supplied interpretations without necessarily changing the labor required to perform them.

Section 7 compares the calibrated organizational reallocation with a task-productivity improvement generating equivalent aggregate reader-capacity savings. To isolate the organizational margin, the productivity counterfactual holds production locations fixed. This restriction is specific to the counterfactual: productivity improvements can themselves induce outsourcing when staffing is indivisible.

The verification mechanism yields four implications for the empirical analysis in Sections 5 and 6:

(i) *Incidence.* Outsourcing is most attractive where local staffing entails large indivisibility penalties, with practice exit potentially amplifying reallocation in thin markets.

(ii) *Client composition.* Conditional on total volume, verification is more valuable to platforms serving many individually thin client relationships, where baseline detection is weaker.

(iii) *Legal feasibility.* Verification can facilitate remote certification only where professional and regulatory requirements permit it; restrictions on remote certification therefore limit this channel.

(iv) *Modality coverage.* Because verification can govern a client relationship rather than a particular task, organizational effects need not be confined to modalities in which AI directly augments reader productivity.

These implications are not individually unique to verification. Generic scale economies can also generate reorganization without a detectable throughput increase, and task-productivity gains may encourage outsourcing. The empirical analysis therefore evaluates their joint consistency with the verification mechanism rather than treating any single pattern as decisive.

4 Data and Empirical Strategy

Our analysis combines four sources of information: FDA clearances identify the availability of reader-facing AI; Medicare claims measure radiologists and interpretations billed across geographic markets; Physician Compare and Doctors & Clinicians data identify practice locations and organizational affiliations; and Medicare Geographic Variation data measure imaging utilization by patients' residence. Together, these sources allow us to distinguish changes in where interpretations are billed and organized from changes in imaging demand. We first describe the data, then construct the model-implied exposure measure and present the empirical strategies.

4.1 AI Technology and Diffusion in Radiology

We construct the technology series from the FDA's list of AI-enabled devices cleared under its radiology panel through March 2026 (U.S. Food and Drug Administration, 2026). Of 1,164 devices, most operate outside the interpretation stage, including acquisition, reconstruction, image processing, and radiotherapy planning. We classify devices by their role in the imaging pipeline and define as *reader-facing* those that analyze image content and return information to the interpreting professional (Appendix A.2). Excluding 30 devices for dental films, which radiologists do not interpret, leaves 249 reader-facing devices.

Of these, 47 percent provide notification or triage, 40 percent second-reader detection or characterization, and 13 percent quantification. None is cleared to issue a final interpretation in place of a licensed radiologist (U.S. Food and Drug Administration, 2026; American College of Radiology, 2025); computer-aided detection, diagnosis and triage devices are classified as aids to the clinician, and Medicare pays for an interpretation only when a physician performs it (Code of Federal Regulations, 2026a,c). Triage and second-reader tools can generate the kind of signal in Assumption 3; quantification acts on the task itself. Because the three diffused together, their separate clocks are not identified (Appendix Table B.3, panel B).

Figure 2 reports cumulative reader-facing clearances by function (panel A) and imaging modality (panel B). Availability expanded rapidly from the late 2010s, with 4 cleared devices in 2013, 18 in 2018, and 194 in 2024.

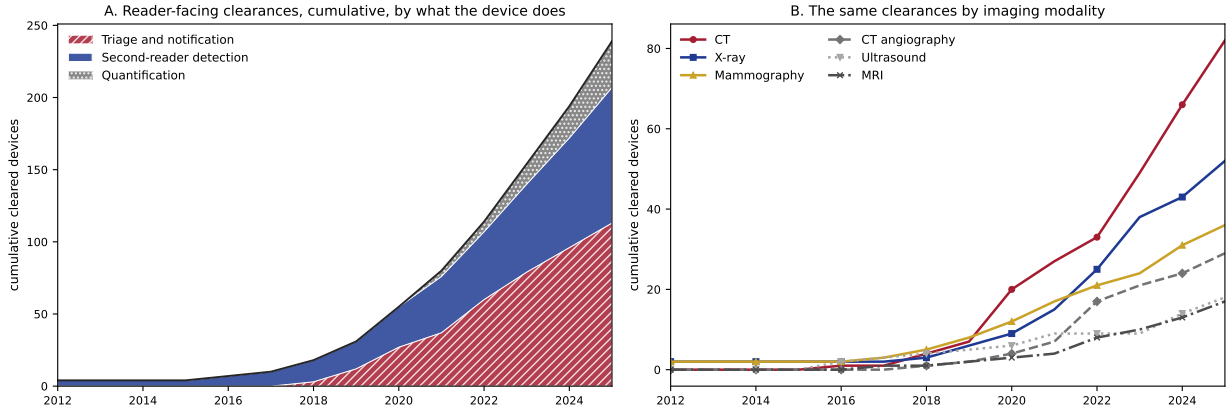


Figure 2: Cumulative reader-facing FDA clearances by function and imaging modality, 2012–2025

NOTE.—Cumulative reader-facing FDA clearances, classified by their role in the imaging pipeline (Appendix A.2). Panel A distinguishes notification and triage, second-reader detection and characterization, and quantification. Panel B reports clearances by imaging modality, used to construct the modality-specific clocks in Equation (11). The aggregate series provides the national clock in Equation (9).
SOURCE.—Authors' classification of the FDA list of AI-enabled medical devices, radiology panel, through March 2026.

Deployed systems assist rather than replace the interpreting radiologist. Commercial applications emphasize detection, quantification, and measurement (van Leeuwen et al., 2021); some re-read examinations against the completed report to identify discrepancies (Topff et al., 2024),

and Pettet et al. (2024) audit the performance of a hemorrhage-detection tool deployed as reporting assistance in a UK teleradiology provider's acute workflow, while randomized evidence examines AI-assisted allocation between single and double reading in breast screening (Lång et al., 2023).

Our baseline measure of AI availability, Clock_t, is the cumulative number of reader-facing devices cleared by year t, entered in logarithms and standardized over 2013–2025. It measures the national technological opportunity set, not adoption by individual firms or markets. Appendix B.3 examines alternative clock transformations and the broader set of radiology-panel clearances. For task-specific tests, we construct modality-specific clocks, Clock_{jt}, from the same devices (Figure 2, panel B); combined with predetermined market modality shares, they yield the task-specific exposure measure in Equation (11).

4.2 Radiology Production and Imaging Demand

Our principal source is the CMS provider-level claims data for 2013–2024 (Centers for Medicare & Medicaid Services, 2026g), supplemented by an archived 2012 release that enters the headcount panel only (National Bureau of Economic Research, 2026). The data identify billing clinicians by National Provider Identifier (NPI), procedures, places of service, and practice addresses. We measure radiologist headcounts and Medicare payments using the by-provider file, and interpretations using radiology procedure codes (HCPCS 70010–79999, plus the Medicare G-codes used for mammography until 2017) from the by-provider-and-service file, excluding the computer-aided-detection add-on codes 77051 and 77052, which were billed beside mammogram reads until 2016 and bundled into the mammography codes from 2017 (6 to 7 percent of services in 2013–16).[4] CMS suppresses procedure cells with fewer than eleven beneficiaries, so interpretation counts are slight undercounts. We identify radiologists by specialty (diagnostic and interventional radiology and nuclear medicine), assign their practice ZIP codes to core-based statistical areas (CBSAs; 2018 delineation, U.S. Office of Management and Budget and U.S. Census Bureau 2018) using the HUD crosswalk (U.S. Department of Housing and Urban Development, 2017), and refer to CBSAs as *markets*. Population comes from Census Bureau CBSA estimates (U.S. Census Bureau, 2025).

Our central outcome, N_{mt}, is the number of distinct radiologists billing Medicare from market m in year t. It measures the geographic location of billed professional work, not where radiologists reside or are employed; Sections 5 and 6 use practice locations, firm affiliations and NPI histories to separate the margins behind it.

We isolate professional interpretations from technical imaging services using place of service. For hospital patients, Medicare pays the physician for the professional component, such as interpretation, while the facility receives payment for the technical component (Centers for Medicare & Medicaid Services, 2025a). We use facility-setting claims, excluding office claims that may bundle professional and technical services. Our baseline production measure, V_{m0}, is the total number

[4] By-provider service totals are unsuitable measures of interpretation volume because they include contrast agents and radiopharmaceuticals billed per unit, which account for 37 percent of reported radiology services in 2013 and 56 percent in 2024.

of facility-setting interpretations in a fixed basket of fifteen high-volume codes covering radiography, CT, MRI, ultrasound, mammography, and nuclear medicine over 2013–14 (codes in Appendix Table A.1); the fixed basket avoids changes in measured volume from procedure-code revisions. Because claims are assigned to providers' practice addresses, V_{m0} measures where interpretations are billed, not where the imaging is performed or where patients reside. We also identify procedures that require the radiologist's physical presence—image-guided biopsies and drainages, paracentesis, thoracentesis, lumbar punctures, and fluoroscopic studies—as a comparison with interpretations that can be supplied remotely (Subsection 6.3).

To test whether reads are done where they are billed, we also use the CMS Physician/Supplier Procedure Summary (Centers for Medicare & Medicaid Services, 2026i), which counts services by payment locality, with ZIP codes assigned to localities from the CMS locality files (Centers for Medicare & Medicaid Services, 2024). For a remote interpretation the claim must carry the ZIP code of the radiologist's location at the time of reading, and that ZIP code sets the locality (Centers for Medicare & Medicaid Services, 2012); for a procedure it carries the place where the service was rendered (Centers for Medicare & Medicaid Services, 2026c). Localities are coarse: the 85 in our panel are 36 whole states, 16 rural remainders of states with several localities, and 33 metropolitan areas. They are set by where the work is done, not by the billing address.

To measure imaging demand independently of billing location, we use CMS Geographic Variation data (Centers for Medicare & Medicaid Services, 2026d), which assign utilization to beneficiaries' places of residence, combining the archived 2007–13 series with the current series for 2014–24 (Appendix Table A.1 documents the alignment across the 2014 classification change). Finally, we construct analogous market-level headcounts and hospital-setting volume for pathologists, and headcounts for family-practice physicians and psychiatrists, as comparison professions. Radiology and pathology exposures correlate 0.35, while their underlying log volumes correlate 0.75, consistent with both services being supplied in the same hospitals.

4.3 Firms, Practice Locations, and AI Adoption

Medicare claims identify billing locations but provide limited information on the organizations connecting radiologists across markets. We recover this structure from five biennial snapshots of Physician Compare and Doctors & Clinicians data for 2017–2025 (Centers for Medicare & Medicaid Services, 2025b), which record physicians' practice locations and group affiliations. We construct two measures of geographic reach: affiliation with a group operating in at least five states, capturing firm reach, and whether a radiologist personally lists practice locations in multiple markets, capturing individual coverage. Both expanded sharply between 2017 and 2025 (Table 1); Subsection 5.4 relates the expansion to the reallocation predicted by Proposition 2.

Linking NPIs across snapshots lets us follow the 2017 cohort of single-location, unaffiliated radiologists through 2024 and the readers who joined platforms between 2019 and 2025, with graduation year (available for 99.9 percent of the cohort) as an age control. We map group identifiers to parent organizations, which identifies readers of the same organization across locations and sepa-

rates documented AI adopters from other multi-state groups, and we build the same organizational measures for other specialties.

Because no comprehensive census of practice-level AI adoption exists, we assemble a separate record of publicly documented events from vendor and firm announcements: 50 dated events involving 24 radiology organizations between 2018 and 2026, 21 of them enterprise-wide deployments. For analyses of dated enterprise adoption, we focus on Radiology Partners with vRad, LucidHealth, and Envision. The record is selective: large firms and vendors have stronger incentives to publicize adoption, while unannounced deployments are unobserved. We do not read missing announcements as evidence of non-adoption, and the main design does not use these events; we use them to describe documented adopters, to study changes within firms around dated deployments, and, in comparisons with other groups (Subsection 5.6 and Appendix C), to separate documented adopters from groups with no documented adoption. We call groups with radiologists in five or more states multi-state groups or platforms; documented adopters are the groups in the adoption record. Survey evidence provides a benchmark: approximately 30 percent of radiologists reported AI use by 2020, concentrated in larger practices (Allen et al., 2021), consistent with evidence that AI and information-technology adoption increases with firm size (McElheran et al., 2024; Lashkari et al., 2024).

4.4 Descriptive Statistics

Table 1 summarizes the main samples and variables. The market panel covers 769 CBSAs[5] and is highly skewed: the median market has eight radiologists billing Medicare in 2013, and 55 percent of markets have fewer than ten, although these account for only about 5 percent of radiologists. Nationally, the number of radiologists billing Medicare increased from 30,708 in 2013 to 35,705 in 2024 (16 percent), while billed imaging interpretations declined from 81.9 million to 79.8 million as the fee-for-service population contracted (Subsection 2.2). The remaining panels summarize hospital interpretation volume, organizational reach, AI availability and documented adoption, and on-site procedures. Appendix Table A.3 reports additional information on the populations and hospitals represented by markets of different sizes.

4.5 Model-Implied Exposure

The model predicts that pooling is especially valuable where local staffing leaves unused capacity. Proposition 2 implies that the reader capacity released by establishments switching to pooled production in market m is

$$\frac{1}{\overline{h}} \sum_{j \in S(\kappa)} \pi(v_j) v_j,$$

[5]Of the 926 CBSAs in the spine, 741 have a billing radiologist in 2013; with the 38 whose first radiologist bills later they form 779 markets, and dropping ten observed in a single year leaves 769, of which 734 have a 2013 headcount (407 with fewer than ten radiologists).

Table 1: Data and summary statistics

A. Markets and the outcome (Medicare claims, 2013 → 2024)	
Markets (core-based statistical areas)	769
Market-years	8,375
Radiologists billing per market, 2013: mean	41.3
median	8
Share of markets with fewer than 10 radiologists, 2013	0.55
share of radiologists in those markets	0.05
Radiologists billing nationally	30,708 → 35,705
Interpretations billed nationally (millions)	81.9 → 79.8
Net change in interpretations, 2013–24 (millions): markets with fewer than 50 radiologists	−3.7
markets with 50 or more radiologists	+9.4
B. Exposure inputs	
Hospital-setting interpretations per market, 2013–14 (thousands): mean	66.9
median	20.3
Population per market, 2015–17 (thousands): mean	402
median	103
Reader-facing AI clearances, cumulative, 2013 → 2018 → 2024	4 → 18 → 194
C. Organization (Physician Compare / Doctors & Clinicians, 2017 → 2025)	
Radiologists listed (distinct NPIs)	32,155 → 35,973
Share in groups operating in five or more states	0.09 → 0.47
Groups with radiologists in five or more states	79 → 448
Share listing locations in more than one market	0.20 → 0.36
Markets listed per radiologist	1.28 → 1.72
Listed radiologists per market, markets with fewer than 10 radiologists	10.8 → 19.8
share of them listing more than one market	0.50 → 0.83
Cohort of single-location, unaffiliated radiologists, 2017	19,382
D. AI adoption record, 2018–2026	
Dated adoption events	50
Adopting organizations	24
Share of listed radiologists at a documented adopter, 2017 → 2025	0.077 → 0.128
E. On-site procedures billed by radiologists (Medicare claims, 2013 → 2024)	
Procedures per market: mean	1,417 → 1,027
median	310 → 136
Share of markets with fewer than 10 radiologists billing none	0.17 → 0.44
Radiologists' share of these procedures nationally	0.80 → 0.72

NOTE.—Panel A is the estimation sample of Equation (9); the market-share rows are among markets with a 2013 headcount, whose 10th–90th percentile range is 1–86 radiologists; the net-change rows give 2024 interpretations minus 2013 interpretations scaled by the change in fee-for-service beneficiaries, summed over the markets in each group; the cohort is followed to 2024. Panel B: raw log volume has a correlation of +0.70 with log population; after residualizing, the exposure of Equation (8) has a correlation of −0.39 with log population. Panel D: the adopter share counts radiologists whose primary organization is one of the 24 in the record, excluding the Strategic Radiology coalition; 10 of the 24 cannot be linked to the directory, so it is a lower bound. Panel E: procedures are attributed to the billing provider's practice address.
SOURCE.—Authors' calculations from the sources listed in Appendix Table A.1.

where v_j is establishment j's interpretation volume and $\pi(v_j)$ its staffing indivisibility penalty. This expression measures potential capacity savings, not local positions displaced. It highlights two determinants of reorganization: the volume of interpretations that can be pooled and the unused capacity associated with local staffing.

Claims measure interpretation volume by billing location but neither establishment-level staffing nor the associated indivisibility penalties. We therefore use predetermined hospital interpretation volume, V_{m0}, to proxy for potentially reorganizable production, and market size to distinguish settings where staffing constraints are more likely to bind. For a given volume, pooling should be more consequential in small markets, where local capacity is harder to share, than in large markets with greater opportunities for local pooling. We estimate the main specifications separately by market size and examine unused reader capacity more directly in Section 5.

We construct exposure as hospital interpretation volume relative to market population. Let $\ln \text{pop}_m$ denote mean log population over 2015–17, and let $\hat{a} + \hat{b} \ln \text{pop}_m$ be the fitted value from regressing $\ln(1 + V_{m0})$ on log population across all markets. Exposure is the standardized residual:

$$Z_m = \frac{e_m - \bar{e}}{\text{sd}(e)}, \qquad e_m \equiv \ln(1 + V_{m0}) - \hat{a} - \hat{b} \ln \text{pop}_m. \tag{8}$$

Thus, one unit of Z_m represents one standard deviation more hospital interpretation volume than expected for a market of that population.

Residualization separates interpretation intensity from market size; Figure 3 illustrates the remaining variation within size classes. The estimating equation also interacts population with the national AI clock. Importantly, V_{m0} is measured in 2013–14, the panel's first years and before the 2018–19 clearance wave, and uses interpretation volume rather than radiologist headcount. This avoids mechanically incorporating the outcome into exposure, as could occur with a headcount-based slack measure such as $N_{m,2013} - V_{m0}/\bar{h}$. Such measures are also difficult to compare across markets because small-market radiologists perform more interpretations per reader than the hub benchmark within the fixed procedure basket. Section 5 examines alternative exposure constructions.

4.6 Empirical Strategy

Baseline market design. Our baseline design interacts predetermined market exposure with the national diffusion of reader-facing AI. Identification does not come from differences in observed adoption across markets: Clock_t measures national technological availability, while Z_m captures predetermined differences in potentially reorganizable interpretation volume. We estimate

$$\ln N_{mt} = \beta \left(Z_m \times \text{Clock}_t\right) + \gamma \left(\ln \text{pop}_m \times \text{Clock}_t\right) + \alpha_m + \delta_t + \varepsilon_{mt}, \tag{9}$$

where α_m and δ_t are market and year fixed effects, and standard errors are clustered by market. We interact population with the clock to account for differential trends in metropolitan concentra-

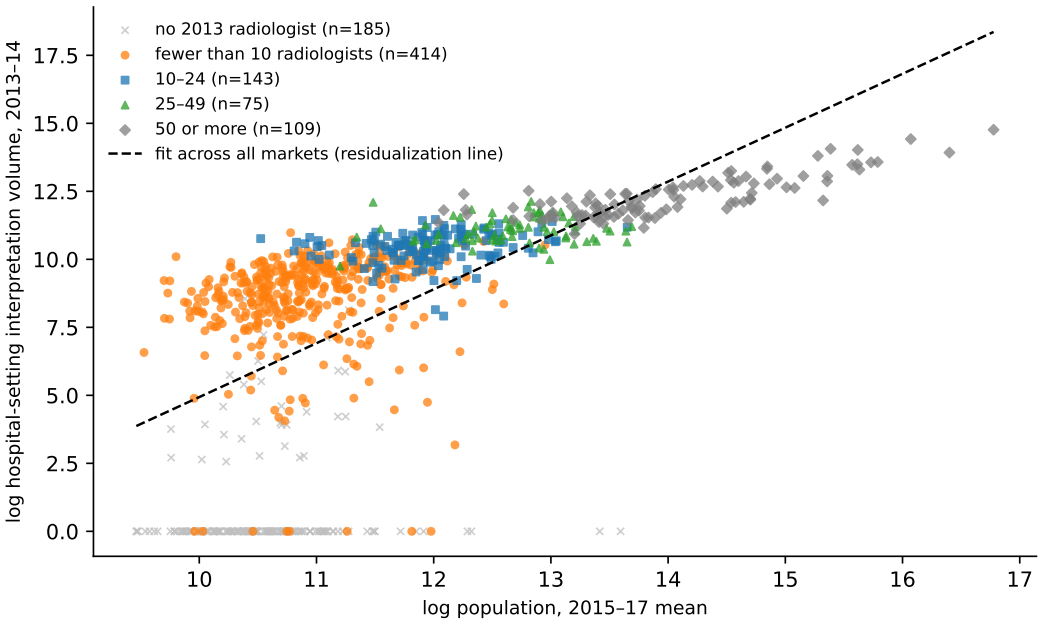


Figure 3: Hospital interpretation volume against population, by market-size class

NOTE.—Each point represents a market, plotting log hospital-setting interpretation volume in 2013–14 against log population. The dashed line is the single fitted relationship of Equation (8) across all 926 markets in the spine, including the 185 with no billing radiologist in 2013 (crosses at zero); exposure is the standardized vertical deviation from this line. Market-size classes are shown for reference and split the 741 spine markets with a billing radiologist in 2013 by headcount.
SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files, 2013–14, and Census population estimates, 2015–17.

tion and report specifications with and without this interaction. Because Z_m and Clock_t are both standardized, β is the change in log headcount per standard deviation of exposure per standard deviation of the clock; the clock rises by 2.5 standard deviations between 2013 and 2024, so cumulative effects by 2024 are about 2.5 times β. We report coefficients as gradients in these units and convert to 2024 magnitudes where stated.

Identification requires that, conditional on market size, markets with different predetermined exposure would not otherwise have experienced differential changes in local radiology billing systematically tracking national AI diffusion. Rather than relying on a discrete break in technology availability, we examine the evolution of the exposure gradient using

$$\ln N_{mt} = \sum_{\tau \neq 2013} \beta_\tau \left(Z_m \times \mathbf{1}[t = \tau]\right) + \sum_{\tau \neq 2013} \gamma_\tau \left(\ln \text{pop}_m \times \mathbf{1}[t = \tau]\right) + \alpha_m + \delta_t + \varepsilon_{mt}. \tag{10}$$

The coefficients β_τ trace the year-specific exposure gradient, allowing us to assess pre-existing differential trends and its subsequent evolution as AI becomes available. Section 5 examines pre-trends, mean reversion, alternative exposure measures and windows, changes in remote-radiology feasibility, extensive-margin outcomes, and comparison professions.

Task-specific exposure. To distinguish task productivity from relationship-level reorganization, let $s_{mj,0}$ denote market m's predetermined share of interpretations in modality j, and Clock_{jt} the corresponding modality-specific AI availability (Figure 2, panel B). We construct

$$\text{TaskAI}_{mt} = \sum_j s_{mj,0} \, \text{Clock}_{jt}, \tag{11}$$

and estimate

$$\ln N_{mt} = \beta \left(Z_m \times \text{Clock}_t \right) + \eta \, \text{TaskAI}_{mt} + \gamma \left(\ln \text{pop}_m \times \text{Clock}_t \right) + \alpha_m + \delta_t + \varepsilon_{mt}. \tag{12}$$

The coefficient η exploits differences in predetermined modality composition. A task-productivity mechanism predicts stronger responses where AI becomes available for locally prevalent modalities, whereas relationship-level verification need not generate modality-specific effects. Because the modality clocks co-move, this test has limited identifying variation; Subsection 5.5 complements it with within-reader evidence.

Within-firm designs. For reader-level outcomes y_{it}, we estimate

$$y_{it} = \phi \left(T_i \times \text{Post}_t \right) + \alpha_i + \delta_t + \varepsilon_{it}, \tag{13}$$

where T_i identifies readers affiliated with a documented enterprise adopter and Post_t indicates the period following adoption. Reader fixed effects absorb time-invariant heterogeneity, and standard errors are clustered by reader. The comparison group comprises readers at multi-state groups without documented adoption; because announcements are selective, these firms cannot be assumed never to have adopted AI.

Using the log of the number of markets or sites per reader, Equation (13) measures changes in geographic coverage per reader. Using the number of Medicare interpretations billed per reader instead allows us to examine changes in individual throughput. A reader-by-task extension, with reader-by-task, reader-by-year, and task-by-year fixed effects, compares the same reader's activity in AI-covered and uncovered tasks around adoption (Subsection 5.5). We also use analogous hospital-level specifications, with hospital and period fixed effects, for turnaround, protocoling, ordering, and subsidy outcomes.

5 The Reorganization of Radiology Production

We examine where radiology production reorganized and whether the resulting patterns accord with the model's four implications (Subsection 3.7): reorganization should be concentrated where staffing indivisibilities make pooling valuable (incidence); verification should be especially valuable to organizations serving many thin client relationships (client composition); remote certification must be legally feasible (legal feasibility); and organizational effects need not follow the imaging

modalities covered by AI (modality coverage). We first establish the incidence and robustness of changes in local billing, then distinguish them from imaging demand, examine organizational responses, and assess task-productivity and institutional alternatives. Section 6 studies the consequences.

5.1 Where the Reorganization Lands

Table 2 reports estimates of Equation (9). The coefficient on exposure × clock is −0.088 ($t = -6.5$): one standard deviation greater predetermined exposure lowers local billing headcount by 8.8 log points per standard deviation of the diffusion clock, which rose by 2.5 standard deviations between 2013 and 2024, so that by 2024 a market one standard deviation more exposed had about a fifth fewer locally billing radiologists. Without the population-by-clock interaction (hereafter, the size control) the coefficient is −0.115. Because the model predicts larger gains from pooling where minimum staffing requirements bind, the split by market size is particularly informative.

Table 2: Exposure and locally billing radiologists, by market size

	2013–14, size-controlled		2013–14, no size control		Markets	*N*
	coef.	(s.e.)	coef.	(s.e.)		
All markets	−0.088	(0.013)	−0.115	(0.011)	769	8,375
Fewer than 10 radiologists	−0.124	(0.019)	−0.100	(0.018)	407	4,266
2–9	−0.125	(0.037)	−0.105	(0.033)	310	3,447
10 or more	0.061	(0.047)	−0.093	(0.016)	327	3,918
25 or more	0.060	(0.053)	−0.038	(0.023)	184	2,208
50 or more	−0.007	(0.057)	−0.031	(0.021)	109	1,308

NOTE.—Dependent variable: log radiologists billing Medicare from the market, 2013–2024. Each cell is the coefficient on exposure × clock from Equation (9), with market and year fixed effects; the size-controlled columns add ln pop×clock. Rows after the first split markets by 2013 billing radiologists; the 10-or-more, 25-or-more and 50-or-more rows are cumulative.
Standard errors in parentheses, clustered by market.
SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files; Census population estimates.

Among markets with fewer than ten radiologists in 2013, the estimated gradient is −0.124 ($t = -6.6$), virtually unchanged at −0.125 ($t = -3.4$) when markets whose counts can move between zero and one are excluded. In markets with ten or more radiologists, the gradient is +0.061 ($t = 1.3$). This concentration in small markets supports the model's incidence prediction, although market size is an imperfect proxy for hospital-level staffing indivisibilities. Indeed, measured unused reader capacity is low in small markets: its median is 0.05 among markets with one to four radiologists and zero among those with five to twenty-four. Single-location radiologists in small markets who still billed in 2024 also performed more interpretations per reader, approximately 4,800 at the median in 2017 versus 2,200 in markets with fifty or more (Table 7). These patterns do not contradict the model: minimum staffing requirements may bind at individual hospitals even when radiologists are heavily utilized across a market; the hospital-level measure of Section 7 is 0.21 in these markets,

against 0.07–0.08 elsewhere (Appendix Table D.2).[6]

The size control matters primarily in larger markets. Among markets with at least ten radiologists, exposure and log population are strongly negatively correlated (−0.92), while population has a pronounced positive gradient (+0.057 per standard deviation, $t = 7.1$), reflecting increasing metropolitan concentration. Controlling for this trend removes the apparent exposure gradient. In smaller markets, the exposure–population correlation is weaker (−0.43), the population gradient is negligible (+0.002, $t = 0.1$), and the exposure estimate changes little.

Reorganization also occurs at the extensive margin. Market-years with no locally billing radiologist account for 9 percent of the balanced panel and occur almost entirely among markets with fewer than ten radiologists in 2013. Their number rises from 336 before 2019 to 497 afterward. In the balanced panel, one standard deviation greater exposure lowers the probability of retaining any locally billing radiologist by 4.6 percentage points per standard deviation of the clock ($t = -4.1$), about 11 points by 2024. Poisson estimates that include the zero cells point the same way but are imprecise (−0.055, s.e. 0.049; small-market interaction −0.100, s.e. 0.060; Appendix Table B.2, panel C). Thus, some markets lose local radiologist billing altogether, although imaging services may continue through interpretations supplied from elsewhere.

The affected markets typically have populations of 50,000–100,000, five to seven billing radiologists, and substantial hospital interpretation volume relative to population. Markets with ten to twenty-four radiologists show no comparable response: their event-study estimates remain close to zero through 2024, with an exposure gradient of −0.05 (s.e. 0.10) and a post-2021 increment of +0.15 (0.09) (Appendix Table B.3, panel A). The decline in local billing is therefore concentrated at the smallest end of the market-size distribution.

5.2 Robustness of the Gradient

We examine whether the exposure gradient reflects mean reversion, alternative dimensions of local production, or specification choices.

Mean reversion is a concern because interpretation volume and radiologist headcount are related: transitory increases in both during the exposure window could generate subsequent declines in headcount without technological change. We first reconstruct exposure using every one-, two-, and three-year window between 2013 and 2017. If transitory fluctuations drove the results, later windows overlapping the outcome period should predict larger subsequent declines. Instead, the gradient attenuates as the exposure window moves later. Using 2015–17 rather than the 2013–14 hold-out period yields estimates of −0.095 versus −0.115 without the size control, and −0.052 versus −0.088 with it (Appendix Table B.1).

Figure 4 provides a second assessment, comparing the event-study estimates from Equation (10) with a calibrated no-technology simulation that reproduces the mechanical relationship between

[6] We approximate unused market-level capacity using $\pi = n\bar{h}/v - 1$ from Subsection 3.1, with 2013 billing radiologists n, interpretations v, and $\bar{h} = 3{,}984$, the 75th percentile of interpretations per billing radiologist. Median values are 0.28 in markets with 25–49 radiologists and 0.53 in markets with 50 or more, where subspecialization, procedures, and part-time work can lower interpretations per reader independently of minimum coverage requirements.

transitory volume and headcount fluctuations (Appendix B.1). The 2012 coefficient, −0.051 (s.e. 0.020), is consistent with a transitory pre-period deviation, while the 2014 coefficient, +0.031 (0.020), lies within the simulated band. Thereafter, the paths diverge: the observed gradient moves outside the simulated 95 percent band in 2019 and reaches −0.22 log points by 2024, while simulated mean reversion remains small. Under this calibration, correction of the initial 2013 deviation accounts for only 5–13 percent of the observed gradient.

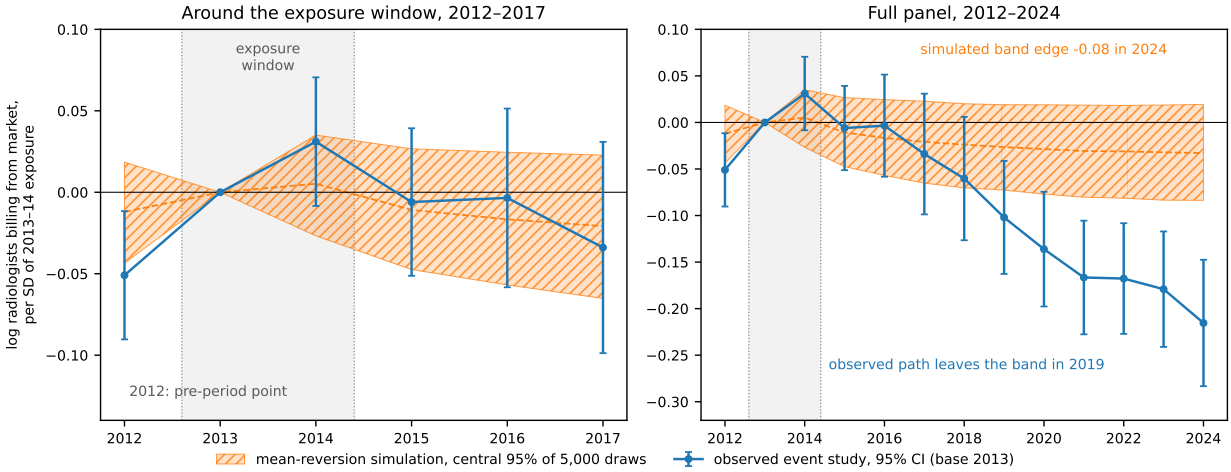


Figure 4: Mean reversion and the evolution of the exposure gradient

NOTE.—Event-study coefficients for 2013–14 exposure (Equation (10); 2013 base year, 2012 pre-period observation), with log population-by-year controls. The simulated 95 percent band comes from 5,000 no-technology replications calibrated to market sizes, idiosyncratic variance, and persistence and estimated using the same procedure (Appendix B.1). Light shading between dotted lines marks the exposure window. Left: 2012–17; right: full panel. The 2012 headcounts use the CY2012 aggregate file (Subsection 4.2).
95 percent confidence intervals; standard errors clustered by market.
SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files, 2013–2024, and the NBER mirror of the 2012 file.

We next examine what economic variation Z_m captures. Appendix Table B.4 shows similar gradients when hospital interpretation volume is residualized on population, entered directly, or expressed per capita (panel A). Controlling for mean volume per billing site, volume at the largest site, or the number of sites also leaves the baseline gradient negative and, if anything, steeper, while these additional measures have coefficients close to zero (panel B). In contrast, exposure measures based on imaging received by market residents produce much smaller gradients: −0.022 (s.e. 0.008) for imaging per beneficiary and −0.034 (0.007) for total beneficiary imaging volume (panel C). The gradient is therefore more closely associated with the location of billed hospital interpretations than with individual-site scale or residents' imaging demand, consistent with reorganization of billing and interpretation relationships.

Finally, the gradient changes little when controlling for Medicare Advantage penetration, interacting pre-period radiologist density with the AI clock, allowing population-specific year effects, or excluding pandemic years individually or jointly. We build density from 2012 headcount, the year before the panel, rather than from 2013 headcount to avoid introducing the mean-reversion

problem discussed above. Weighting markets by their 2013 radiologist headcount, however, yields a much smaller and imprecise coefficient (−0.015, s.e. 0.044), showing that the baseline gradient characterizes the average market rather than the average radiologist. Estimates remain negative in metropolitan (−0.072, s.e. 0.032) and micropolitan (−0.093, s.e. 0.016) markets. Alternative AI-clock transformations retain the same sign, and randomization inference using 499 permutations places the observed estimate in the extreme tail of the permutation distribution (Appendix B.3).

5.3 Billing Moved; Imaging Demand Did Not

Our baseline outcome measures where radiologists bill, not where they work or where patients are imaged. Table 3 separates the margins.

Imaging demand has no exposure gradient: imaging received per beneficiary moves by −0.001 (s.e. 0.001), and payment per service is unchanged (panel C). Affiliations and billing diverge: in the directory years the gradient in locally billing radiologists is −0.091 ($t = -4.4$), against +0.027 ($t = 1.2$) for all radiologists listing a location in the market and −0.090 ($t = -5.8$) for those listing only that market (panels A and B).[7] In markets with fewer than ten radiologists, listed radiologists per market rise from 10.8 to 19.8 between 2017 and 2025 and the share listing several markets from 50 to 83 percent (Table 1).

Volume moved more than people. Hospital-setting interpretations billed from a market fall by 0.34 log points per standard deviation of exposure ($t = -7.0$), four times the headcount gradient; across all imaging codes, facility-setting interpretations fall by 0.27 (0.44 in small markets) while office-setting interpretations do not (+0.17, $t = 2.6$) (panel A). In levels, facility interpretations billed from small markets fall by about half between 2013 and 2024 (to 0.49 of their 2013 level), against 6 percent in the largest markets (Appendix Table D.3), and relative to their own fee-for-service demand, markets with fewer than fifty radiologists bill 3.7 million fewer interpretations in 2024 than in 2013 while markets with fifty or more bill 9.4 million more (Table 1). Levels also carry trends that predate AI, so identification rests on the clock-timed gradient; but the pattern is a reallocation of hospital interpretation billing, not a contraction of local imaging demand.

Who left, and where did they go? Assigning each radiologist permanently to the market of first billing and counting the NPI wherever it bills gives a measure that moves only with entry and exit. Its gradient is −0.029 ($t = -3.5$) overall and −0.053 ($t = -4.9$) in small markets, against −0.088 and −0.124 on current addresses (Appendix Table C.3): in small markets, reduced Medicare participation accounts for about two-fifths of the decline and relocated billing for the rest. Among small-market radiologists who were not platform-affiliated in 2019 and billed elsewhere by 2024, three-quarters no longer list the origin market and two-thirds bill from a market with fifty or more radiologists; their billed services fell about as much as those of radiologists who stayed (−0.38 against −0.35 log points). Relocation is not a paper transfer.

Reads physically done in a locality fall with exposure as much as reads billed from it. Across

[7] Estimated on the directory snapshot years, to match the listing rows; claims end in 2024, so the billing row uses 2017–2023.

Table 3: What moved and what did not: exposure gradients, margin by margin

Outcome	coef.	(s.e.)	N
A. Billed locally (the margin that moves)			
Radiologists billing from the market (annual panel, Table 2)	−0.088	(0.013)	8,375
Radiologists billing from the market (directory years)	−0.091	(0.021)	2,706
Hospital interpretation volume billed from the market (exposure basket)	−0.339	(0.048)	8,315
all imaging codes, facility setting	−0.274	(0.043)	8,212
facility setting, markets with fewer than 10 radiologists	−0.436	(0.058)	4,129
office setting, all markets	0.170	(0.066)	6,743
office setting, markets with fewer than 10 radiologists	0.141	(0.091)	2,783
Total Medicare payments billed from the market	−0.169	(0.037)	8,375
B. People (conserved, recomposed)			
All radiologists listing a location in the market	0.027	(0.022)	4,161
of whom: listing *only* this market	−0.090	(0.015)	4,161
C. Output and prices (conserved)			
Imaging received per resident beneficiary	−0.001	(0.001)	10,841
Payment per service	−0.004	(0.017)	8,375
D. Where reads are physically done: payment localities, 2016–2024			
Facility-setting interpretations read in the locality	−0.101	(0.017)	765
The same interpretations billed from the locality	−0.094	(0.020)	765
Difference, read minus billed	−0.007	(0.013)	765

NOTE.—Each row is the coefficient on exposure × clock from Equation (9) with the size control. Annual rows use the estimation sample of Table 2; the hospital-volume row keeps market-years with a positive volume. Directory rows use the biennial Physician Compare snapshots over all markets, 2017–2023 for the billing row and all five for the listing rows; the exclusively local count enters as ln(count + 0.5). The beneficiary row uses Medicare's Geographic Variation files, 2013–2024, which record imaging by the patient's residence. The directory rows cover the 871 spine markets with a listed radiologist in at least one snapshot. Panel D: log volume in the 85 CMS payment localities (2016 definitions) on locality exposure × clock with locality and year fixed effects, where locality exposure is the 2016 volume-weighted mean of market exposure; weighted by 2016 billed volume. Read in the locality: Physician/Supplier Procedure Summary, where the claim's ZIP code is the radiologist's location at the time of reading; billed: the claims files at the practice ZIP.
Standard errors in parentheses, clustered by market (panel D: by locality).
SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files; the Physician Compare and Doctors & Clinicians files; the Medicare Geographic Variation files; the Physician/Supplier Procedure Summary; the CMS ZIP-to-locality files; Census population estimates.

the 85 payment localities, 2016–2024, facility-setting interpretations performed in the locality have an exposure gradient of −0.101 (s.e. 0.017), against −0.094 (0.020) for interpretations billed from it; the difference is −0.007 (0.013) (Table 3, panel D). In the rural parts of states with several localities, reads performed fell to 0.80 of their 2016 level and reads billed to 0.72. The reads left these localities, not only the billing addresses; localities are too coarse, however, to rule out moves within them.

None of this requires fewer radiologists in aggregate: the profession grew 16 percent while locally billed work in exposed markets moved elsewhere (Subsection 6.1), and Subsection 5.5 shows the movement does not follow AI-covered tasks or coincide with more reads per reader.

5.4 Platforms, Adoption, and Geographic Reach

The decline in local billing coincided with rapid expansion of multi-market radiology organizations: between 2017 and 2025, the share of radiologists affiliated with groups operating in five or more states rose from 9 to 47 percent, and the number of such groups from 79 to 448. The model predicts that verification is most valuable to organizations that can spread its fixed cost across many client relationships, especially thin ones, and that pooling expands those organizations' geographic reach.

Documented adoption fits the scale part of the first prediction. Most of the 24 organizations in our record are national or multi-state groups—Radiology Partners, SimonMed, RadNet, LucidHealth, vRad, Envision, Akumin, Rayus—or members of Strategic Radiology, through which independent practices share fixed costs (Strategic Radiology, 2026); the rest are regional groups and teleradiology firms, and no single-site department appears. Public announcements are selective, so undocumented adoption cannot be treated as non-adoption. Two observations point the same way: an American College of Radiology survey finds greater AI use in larger practices (Allen et al., 2021), and hospitals in markets with fewer than ten radiologists switch to platform coverage more often than those in the largest markets (13.6 versus 9.9 percent), particularly when they have few radiologists of their own; at the switch, the hospital's locally located radiologists fall by about a quarter and its remotely located ones rise by as much (Appendix C.1).

Table 4 tests the second. At Radiology Partners, markets listed per reader rose from 1.41 in 2017 to 2.41 in 2025, while changing little at multi-state groups without documented adoption (panel A); Equation (13) puts the relative increase at 17 log points in markets and 29 in practice sites per reader (panel B). Panel C guards against directory conventions by following radiologists unaffiliated with any platform in 2019: those who joined Radiology Partners went from 1.52 to 2.54 listed markets, against 1.62 to 1.76 for joiners of other multi-state groups and 1.36 to 1.43 for those who stayed unaffiliated. Greater reach is not a mechanical consequence of joining a large group, though it is not a causal effect of AI either: adoption, growth and client acquisition are jointly determined.

Greater reach is not accompanied by higher throughput. At Radiology Partners, Medicare interpretations per reader fall by 4.6 log points after 2021, and at LucidHealth by 12.5 after 2020 (panel B). Adopters expanded their coverage without reading more per reader, which is hard to reconcile with an explanation running only through reading speed, although gains outside Medicare remain possible.

Finally, Table 5 places radiology among other specialties. In 2025, 47 percent of radiologists belong to groups listing radiology in five or more states, against about 5 percent in the next specialty (neurology); on the broader any-organization definition the share rose from 10 to 48 percent in radiology, against 2 to 10 in emergency medicine and 1 to 9 in anesthesiology. Multi-state consolidation is not unique to radiology, but its scale and geographic form are. Generic scale economies build large organizations without moving where services are supplied; radiology combines remotely deliverable interpretations with technologies that can generate information about professional output. The comparison cannot separate verification from other technologies of remote organization;

Table 4: Platforms: coverage per reader, the reader-level estimates, and the joiner comparison

A. Markets covered per listed reader, by organization type

Primary group	2017	2019	2021	2023	2025
Radiology Partners (adopter; enterprise agreement 2021)	1.41	1.71	2.15	2.37	2.41
RadNet (adopter)	1.85	2.00	1.72	2.21	2.42
vRad (own tools before 2020)	1.77	1.86	1.95	1.92	—
Other multi-state groups (no documented adoption)	1.92	1.91	1.74	1.90	1.73
Single-market practices	1.19	1.37	1.37	1.43	1.43

B. Adopter readers against readers at other multi-state groups: adopter × post, reader and year fixed effects

Outcome, per reader	coef.	(s.e.)	Readers	N
Log markets listed, Radiology Partners	0.174	(0.020)	4,729	14,172
Log practice sites listed, Radiology Partners	0.293	(0.028)	4,729	14,237
Log Medicare interpretations, Radiology Partners, post-2021	−0.046	(0.017)		208,247
Log Medicare payments, Radiology Partners, post-2021	−0.031	(0.016)		208,247
Log Medicare interpretations, LucidHealth, post-2020	−0.125	(0.059)		195,316

C. Readers with no platform affiliation in 2019, markets listed per reader, by 2025 destination

2025 destination	2019	2025	Change	Readers
Joined Radiology Partners	1.52	2.54	+1.02	949
Joined RadNet	1.28	2.08	+0.80	194
Joined a multi-state group with no documented adoption	1.62	1.76	+0.14	1,958
Remained unaffiliated	1.36	1.43	+0.07	13,656
Difference, Radiology Partners joiners minus other multi-state joiners			+0.88 (0.068)	

NOTE.—Panel A: mean number of markets in which a listed radiologist reports a practice location, by the organization type of the radiologist's primary group; vRad is absent from the 2025 file. Panel B: Equation (13) with reader and year fixed effects; the listing rows use the five directory snapshots, the Medicare rows annual claims 2017–2024; the comparison group is readers at multi-state groups with no documented adoption for the listing rows and readers at independent, academic, local and Strategic Radiology groups for the Medicare rows. Panel C: readers who belonged to no platform in 2019, by where they were in 2025; the last row is the difference in the 2019–2025 change between joiners of Radiology Partners and of non-adopting multi-state groups. Post = 2021 for Radiology Partners (enterprise agreement) and 2020 for LucidHealth (deployment).
Standard errors in parentheses, clustered by reader (panel B); for the last row of panel C, the standard error of a two-sample difference in means with unequal variances.
SOURCES.—Authors' calculations from the Physician Compare and Doctors & Clinicians files; the Medicare Physician & Other Practitioners files; the authors' record of documented AI adoption.

the next subsections take those up.

5.5 Task Productivity and Firm Reorganization

The reorganization (implication iv) could also come from task productivity: if AI raised interpretations per radiologist, it could relax staffing indivisibilities and move firm boundaries without improving verification, the channel the model carries through $a(\kappa)$ (Subsection 3.7). Regulation rules out substitution—no reader-facing system studied here can certify an interpretation (Subsection 4.1)—but augmentation remains possible, and a mechanism running through it should raise measured interpretations per reader, above all in the tasks the tools cover. Table 6 reports the tests with 95 percent confidence intervals, so that the magnitudes the data exclude are visible alongside the absence of detectable effects.

Table 5: Multi-state organization of practice by specialty, 2017 and 2025

Specialty	NPIs (2025)	Group lists this specialty in ⩾ 5 states (%)		Group lists any specialty in ⩾ 5 states (%)		AI devices cleared by 2024
		2017	2025	2017	2025	
Radiology (DR, IR, NM)	35,973	9.1	46.6	9.6	48.1	841
Pathology	12,220	0.7	0.9	1.6	12.1	5
Neurology	17,156	0.4	5.2	1.3	16.4	46
Emergency medicine	39,394	1.1	2.4	1.8	9.6	—
Anesthesiology	41,667	0.4	0.9	1.3	9.4	18
Eight other specialties, range	308,136	0.1–1.0	0.3–3.7	0.3–1.6	4.6–17.7	2–110

NOTE.—Percent of physicians with the stated primary specialty, from the Physician Compare (2017) and Doctors & Clinicians (2025) files. Group lists this specialty: the physician's primary organization lists that specialty in five or more states. Group lists any specialty: it lists any specialty in five or more states, which catches multi-specialty staffing companies. Last row: the range across hospitalists (no 2017 code), dermatology, ophthalmology, psychiatry, internal medicine, family practice, cardiology and orthopedic surgery; its first column is their total. NPI: National Provider Identifier; DR, IR, NM: diagnostic, interventional, nuclear.
SOURCES.—Authors' calculations from the Physician Compare and Doctors & Clinicians files; the FDA list of AI-enabled medical devices, by lead review panel, through 2024.

Table 6: Task margins: task mix, firms, and outcomes

Margin	Estimate	(s.e.)	95% CI	N
A. Task mix: AI-covered against uncovered tasks				
Within readers at Radiology Partners, post-2021	−0.044	(0.023)	[−0.090, 0.001]	1,355,172
Across markets, by exposure	0.041	(0.059)	[−0.076, 0.157]	61,285
Modality shift-share on headcount, equation (12)	−0.054	(0.113)	[−0.276, 0.168]	8,210
volume exposure × clock in the same regression	−0.086	(0.016)	[−0.116, −0.055]	8,210
B. Firms				
Markets covered in 2017 by the AI-using teleradiology firm, exposure gradient	0.012	(0.020)	[−0.027, 0.051]	8,375
C. Outcomes (percentage points)				
Emergency head CT read within 45 minutes, hospitals covered by an adopter in 2019	−1.06	(0.95)	[−2.92, 0.80]	12,707

NOTE.—Panel A: reader × task triple difference on claims 2017–2024 (first row); market × task triple difference on the 2013–2024 panel with task-by-year effects by 2013 market size (second row; with common task-by-year effects −0.084 (0.032)); the modality shift-share with market and year fixed effects (third and fourth rows). Panel B: market-level difference-in-differences on the 2013–2024 panel. Panel C: hospital-level difference-in-differences on reporting periods 2018–2025; switching hospitals are in Table 8. Panels B and C are detailed in Appendix C.2.
Standard errors in parentheses and 95 percent confidence intervals in brackets. Panel A: clustered by reader (first row) and by market (other rows); panel B: clustered by market; panel C: clustered by hospital, N hospital-periods.
SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files; CMS Hospital Compare; the Physician Compare and Doctors & Clinicians files; the authors' record of documented AI adoption.

Aggregate throughput shows no upward break. Interpretations per billing radiologist rose 0.4 percent a year in 2013–16 and fell 0.3 percent a year across 2016–19, the window containing the clearance wave; they fell 6.4 percent a year in the pandemic years 2019–21 and a further 1.7 percent a year in 2021–24, after the enterprise deployments (Figure 1). Interpretations per fee-for-service beneficiary rose 0.4, 1.3 and 2.7 percent a year over the same non-pandemic windows, so the post-pandemic shortfall in reads per radiologist reflects fewer fee-for-service patients per radiologist (4.4 percent a year fewer after 2021), not fewer reads per patient. Neither series rises with AI diffusion.

Within the largest documented adopter, covered tasks do not expand relative to uncovered ones. Radiology Partners deployed at least one tool to about 2,600 of its 3,600 radiologists from 2021 (Radiology Partners, 2021, 2023), covering head CT, chest radiography and mammography among others. Relative to uncovered tasks and to readers at comparison groups, covered-task interpretations fall by 0.044 log points (confidence interval −0.090 to +0.001), so the data exclude a meaningful relative increase (panel A); interpretations fall in covered and uncovered radiography and CT alike, and mammography does not rise despite a relevant tool. Across all tasks, interpretations per reader at the firm fall by 0.046 log points while its geographic reach expands (Table 4).

Across markets, local billing does not follow modality-specific AI availability. In Equation (12) the task-specific exposure has a coefficient of −0.05 ($t = -0.5$) beside −0.086 ($t = -5.5$) for volume exposure; the task-specific estimate is imprecise because the modality clocks co-move. Covered reads do not decline more than uncovered reads once task-by-year effects are allowed to differ by market size (panel A). Markets served in 2017 by an AI-using teleradiology firm, which used AI to speed reading rather than to contract for it, lose no more locally billing radiologists than other markets (panel B), emergency head-CT turnaround does not improve at hospitals covered by adopters (panel C), and a $1,040-per-use stroke-AI subsidy (Centers for Medicare & Medicaid Services, 2020b) produces no on–off response in thrombectomy volume or turnaround (Appendix C.2).

Two limitations apply. We cannot say how employment or firm boundaries would respond if AI could certify interpretations, and we do not observe diagnostic accuracy, missed findings, patient health or the effort per read; the absence of measured throughput gains is not the absence of efficiency or quality gains. Nor do the results require $a'(\kappa) = 0$: productivity gains may exist along unmeasured margins and may themselves facilitate pooling. What they rule out is an explanation running solely through higher measured reading throughput, which cannot account for the geographic and organizational evidence. We next ask whether the reorganization follows the institutional restrictions the verification mechanism predicts.

5.6 Alternative Explanations and Institutional Falsifications

Several alternatives could explain the reorganization—local economic decline, generic physician-practice consolidation, including private-equity roll-ups, or institutional changes that eased remote interpretation independently of AI. The pathology comparison tests the market-level alternatives at once; the licensure, ownership and entry tests then take the radiology-specific ones in turn.

Local decline first. If high-exposure markets were losing hospitals, reduced local radiologist

billing might reflect contracting demand for hospital interpretations. Yet between 2019 and 2026, neither the number of acute-care and critical-access hospitals nor the probability that a hospital disappears from CMS records has a meaningful exposure gradient.[8] Employment in image acquisition does not fall with exposure either: radiologic technologists rise slightly (+0.021, $t = 2.6$) and sonographers, nuclear-medicine technologists, family-practice physicians and psychiatrists show no gradient (Appendix Table B.5).[9]

Pathology offers a more demanding comparison: pathologists interpret specimens for many of the same hospitals, bill the same payer, and went through the same consolidation of physician practices. Remote certification differed, however: each site where a pathologist signs out cases needs its own laboratory certificate, so until March 2020 pathologists could not sign out from an uncertified remote site,[10] and the legal-feasibility prediction implies less scope for reorganizing pathology before 2020.

The patterns differ sharply. The share of pathologists affiliated with groups operating in five or more states rises only from 0.7 to 0.9 percent between 2017 and 2025. The exposure gradient in local pathologist billing is +0.009 (s.e. 0.018), against −0.088 for radiology; it is −0.004 (0.021) in markets with fewer than ten pathologists and +0.054 (0.041) in larger markets. Nor is there an extensive-margin decline: the gradient in the probability of retaining any locally billing professional is −4.6 percentage points for radiology and +5.3 for pathology ($t = 3.2$) (Appendix Table B.5, pathology rows). Year by year, the pathology gradient never moves: relative to 2019, every coefficient of its event study for 2013–2024 lies within 0.02 of zero.

Relaxing remote sign-out in March 2020 did not produce a radiology-like adjustment: the five-state share stayed below 1 percent while radiology's rose from 9 to 47 percent (Table 5), and the pathology gradient does not become detectably more negative.[11] Legal permission appears insufficient without complementary investment, notably digitized slides. Because the two specialties share hospitals, payer and market-level shocks, the pathology null rules out the market-level explanations listed above at once and places the radiology gradient in what distinguishes the two specialties: reader-facing AI (841 AI devices cleared on the radiology panel through 2024 against 5 for pathology, Table 5), and, until 2020, the legal scope for remote certification. It cannot by itself separate verification from other radiology-specific channels; the remaining tests address those.

Interstate licensure next. States joined the Interstate Medical Licensure Compact at different

[8] CMS Hospital General Information snapshots, 2019–2026: the exposure gradient in the log of one plus hospitals is +0.001 (s.e. 0.008); among hospitals listed in 2019, 8.1 percent disappear by 2026, with a gradient of +0.013 (0.017). Conversions to rural emergency hospitals count as exits (Centers for Medicare & Medicaid Services, 2026j).

[9] The technologist gradient reflects growth before 2017; in 2017–2025 it is +0.007 (s.e. 0.012). The BLS occupational employment survey (U.S. Bureau of Labor Statistics, 2026) covers metropolitan areas only, 83 of the 431 CBSAs with fewer than ten radiologists; among the 77 with data the gradient is +0.008 (−0.021 to +0.037).

[10] A separate certificate per location: 42 CFR 493.35, 493.43 and 493.55 (Code of Federal Regulations, 2026b), waived for remote review in the public health emergency (Centers for Medicare & Medicaid Services, 2020a) and since kept for digital slides (Centers for Medicare & Medicaid Services, 2023); the FDA authorized whole-slide imaging for primary diagnosis in April 2017 (U.S. Food and Drug Administration, 2017).

[11] Post-2020 change in the slope of the pathology gradient: −0.005 per year (−0.026 to +0.016) overall and −0.006 (−0.034 to +0.023) in markets with fewer than ten pathologists; both intervals include radiology's pre-2020 small-market trend of −0.032.

dates, lowering one cost of supplying services across state lines. Compact entry is associated with a change in local radiologist billing of +0.016 ($t = 0.9$), about +0.03 in small markets; event-study estimates stay near zero from four years before to five years after entry, and entry changes neither multi-state affiliation nor the share of radiologists listing several markets (Appendix C.5). The compact × exposure interaction is −0.036 (s.e. 0.029). Teleradiologists working across state lines already needed licenses in each receiving state, and Texas, Florida, California and New York had not joined the compact by 2020 (Hanna et al., 2020).

Private-equity ownership does not account for the estimated exposure gradient. Adding each market's share of radiologists in private-equity-backed groups and its interaction with exposure leaves the gradient at −0.087 (s.e. 0.013), while the ownership interaction is −0.024 (0.095) (Appendix Table B.2, panel A). This does not rule out a complementary role for private-equity financing in the expansion of multi-market platforms.

Credentialing and other technologies of remote organization travel with platform expansion, so we compare entry by groups documented as using interpretive AI, from their first deployment, with entry by other multi-state groups. Entry alone has no detectable association with local billing headcount for either group. In small markets, however, the exposure gradient is larger where documented adopters enter: the adopter-entry interaction is −0.24 per standard deviation ($t = -2.1$), against +0.015 ($t = 0.2$) for other multi-state entrants (Appendix C.6), although the difference between the two is not statistically significant and the pattern reverses sign in markets with ten to forty-nine radiologists. Because platform entry is endogenous and documented AI adoption is selectively observed, these estimates do not identify the causal effect of AI and are suggestive only.

6 What the Reorganization Did

Section 5 documented a geographic reallocation of radiology production consistent with verification facilitating the remote organization of professional work, rather than an explanation based solely on increased reading throughput. We now examine its consequences for aggregate and local employment, radiologists initially billing from affected markets, and the hospitals they serve.

6.1 Employment, Net and Gross

The reorganization did not coincide with a national contraction in Medicare-billing radiologists. Their number rose 16 percent, from 30,708 in 2013 to 35,705 in 2024, even as local billing declined in exposed markets—consistent with Proposition 2: pooling can eliminate local positions while aggregate employment grows if demand for radiologist services outpaces the capacity that reorganization releases.

The local loss is nonetheless concrete. Had above-average-exposure markets experienced average exposure, about 840 more radiologists (95 percent confidence interval 580–1,120) would have billed

from them in 2024, about 440 of them in markets with fewer than twenty-five radiologists.[12] Not all of these positions vanished from the profession: as Subsection 5.4 showed, reading moved into multi-market organizations.

Whether the loss matters depends on the denominator. The 431 CBSAs with fewer than ten radiologists in 2013 (407 in the estimation sample) hold about 5 percent of the profession but, together with markets lacking a locally billing radiologist and rural areas outside CBSAs, two in five hospitals and three in four critical-access hospitals, serving about one in six Americans (Appendix Table A.3). The reorganization is modest relative to national radiologist employment and consequential for the hospitals and communities where local billing is thinnest.

One thing the expansion did not do is raise local concentration of listed affiliations. Grouping the radiologists listing a location in each market by parent organization, the Herfindahl index declines on average between 2017 and 2025, in small and large markets alike, with no exposure gradient.[13] The index measures listed affiliations, not shares of interpretations, revenues or hospital contracts, so it does not establish that concentration in service provision fell.

6.2 What Happened to Local Radiologists

Declining local billing can reflect three individual adjustments: ceasing to bill Medicare, billing from another market, or staying and billing less. Table 7 follows them for the 19,382 radiologists who in 2017 billed from a single location and belonged to no multi-state group.

Table 7: Where the 2017 cohort of single-location, unaffiliated radiologists was in 2024

		Status in 2024 (%)			Stayers' interpretations			Five-state
2013 market size	N	Exited	Moved	To hub	2017	2024	Δ ln	group, 2025 (%)
Fewer than 10 radiologists	742	30.3	10.8	60.0	4,758	2,931	−0.41	34.8
10–49 (pooled)	2,621	25.1	9.0	72.9	4,076	2,731	−0.36	44.7
50 or more	16,014	20.9	7.6	80.5	2,164	1,688	−0.23	39.7
All	19,382	21.8	7.9	78.3	2,433	1,826	−0.25	40.2

NOTE.—Radiologists who in 2017 billed Medicare from a single location and belonged to no group operating in five or more states, by 2013 market size (10–24 and 25–49 pooled; five members without a 2013 market class are in the total only). Exited: no longer billing Medicare in 2024; moved: billing from another market; to hub: share of movers whose destination has 50 or more radiologists; stayers' columns: medians of annual interpretations and the median within-radiologist change in logs; last column: share of those still billing in 2024 who are affiliated with a five-state group in the 2025 Doctors & Clinicians file.
SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files; the Physician Compare and Doctors & Clinicians files.

All three are larger in the smallest markets. Among radiologists starting in markets with fewer

[12] For each above-average-exposure market, counterfactual 2024 headcount equals actual headcount multiplied by $\exp(-\hat{\beta} z_m \Delta\text{Clock})$, where $\hat{\beta}$ is estimated from Equation (9) and ΔClock is the change from 2013 to 2024. Because year fixed effects absorb common changes, this calculation measures differences in local billing relative to markets with average exposure, not the number of radiologists leaving the profession.

[13] Radiologists listing a location are grouped by parent organization in each directory snapshot; those listing no organization are treated as individual groups. The mean index declines from 0.59 to 0.52 overall, from 0.73 to 0.66 in markets with fewer than ten radiologists, and from 0.42 to 0.36 elsewhere. Its exposure gradient is +0.013 (s.e. 0.011) overall and +0.013 (0.012) in small markets; the gradient in the largest group's share is +0.009 (0.010).

than ten radiologists, 30 percent had stopped billing Medicare by 2024 against 21 percent in the largest markets; 11 percent billed from another market, three in five of them from a hub (a market with fifty or more radiologists); and the stayers' median annual interpretations fell by 0.41 log points against 0.23. The pattern is in line with the calibrated model of Section 7, in which local units holding 21 percent of small-market reader positions fall below the one-reader viability threshold, 8 points of them because of reallocation and the rest on demand alone.

Age explains part of the exit gap, not all of it. Small-market radiologists graduated earlier; conditioning on graduation year reduces the 9.4-point gap between the smallest and largest markets to 4–5.5 points ($t = 2.9$ to 3.6),[14] and giving initial age composition its own time trend leaves the market-level gradient at -0.092 against -0.086 (Appendix Table B.2).

A market-level decomposition points the same way. The gradient in the stock of surviving 2013 incumbents is -0.097 ($t = -7.3$), but exit rates show no differential gradient across tenure classes, and neither entry rates nor the stock of entrants fall with exposure (Appendix Table C.3, panel B). The decline came from incumbents, not from fewer entrants.

None of this measures earnings. Claims record Medicare billing, not compensation, and omit other payers; we cannot say how the reorganization affected the pay of those who moved to hubs, joined platforms, or stayed. What is established is the adjustment in billing participation, location and volume.

6.3 Local Procedures and Hospital Care

Procedures such as biopsies, drainages and lumbar punctures cannot be supplied remotely, so a thinner local radiologist presence could reach services that AI does not touch.

Panel A of Table 8 applies Equation (9) to radiologist-performed procedures in markets with fewer than ten radiologists. Procedures billed by radiologists fall by 0.26 log points per standard deviation of exposure, about twice the headcount gradient, and the probability of any radiologist-billed image-guided procedure by 6 percentage points. Other physicians, nurse practitioners and physician assistants do not fill the gap, so procedures billed by anyone fall too (Appendix C.4).

Claims place a procedure at the radiologist's billing address, not where it was done. Payment-locality counts, which record the place of service, attribute about half of the exposure-related decline to procedures performed locally and the rest to where they are billed; fluoroscopic studies fall mostly where they are performed, image-guided procedures mostly where they are billed (Appendix Table C.4, panel C). Reassigning procedures to small markets where the billing radiologist lists a practice location points the same way: this deliberately generous correction removes about three-fifths of the gradient in the $\ln(1+n)$ specification of panel D (-0.19, s.e. 0.07, against -0.46 before reassignment; Appendix Table C.4) but not all of it.[15]

[14] Among the 19,367 cohort members with a graduation year: graduation-year bins give 5.5 points (s.e. 1.5), a quadratic 4.3 (1.5).

[15] Procedures billed by radiologists with a claims address outside a small market are assigned to it if they list a practice location there (claims years 2017, 2019, 2021, 2023 and 2024; split equally across qualifying markets). Adding hospital affiliations, at most three per radiologist, reassigns two to three times as many and eliminates the gradient

Table 8: On-site procedures in small markets, and hospitals after a switch to platform coverage

A. Procedures that need a radiologist present, markets with fewer than 10 radiologists	Log count			Probability of any		
	coef.	(s.e.)	N	coef.	(s.e.)	N
Billed by radiologists: all on-site procedures	−0.260	(0.076)	3,410	−0.048	(0.016)	4,848
image-guided procedures	−0.230	(0.097)	2,457	−0.061	(0.014)	4,848
fluoroscopic studies	−0.249	(0.096)	3,119	−0.064	(0.016)	4,848
Billed by other physicians and practitioners	0.007	(0.061)	1,668	−0.034	(0.013)	4,848
Billed by anyone	−0.192	(0.056)	3,727	−0.040	(0.013)	4,848

B. Hospitals that switched to an adopting platform: change after the switch (percentage points)	coef.	(s.e.)	N	Switching hospitals
Emergency head CT read within 45 minutes	−3.76	(1.61)	11,831	193
Abdominal CT scanned with and without contrast	0.05	(0.20)	27,868	302
Placebo: MRI of the lumbar spine for low back pain	−0.39	(0.66)	7,535	104
Placebo: cardiac imaging before low-risk surgery	−0.08	(0.13)	14,468	176

NOTE.—Panel A: coefficients on exposure × clock from Equation (9) with the size control, 2013–2024; the log column uses market-years with a positive count, the probability column the balanced panel; procedures as in Appendix Table C.4, at the billing provider's practice address. Panel B: switching hospitals against never-covered hospitals, hospital and reporting-period fixed effects, CMS Hospital Compare outpatient measures 2018–2025; the head-CT share is a timeliness target (higher is better) and its decline began two to three years before the switch (Appendix C.2); the other three are shares of imaging ordered by referring clinicians (lower is better).
Standard errors in parentheses, clustered by market (panel A) and by hospital (panel B); N: market-years (panel A), hospital-periods (panel B).
SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files; CMS Hospital Compare; the Physician Compare and Doctors & Clinicians files; the authors' record of documented AI adoption.

Whether patients traveled or went without, claims cannot say. Larger markets within 150 kilometers of exposed small markets bill more procedures (+0.028, s.e. 0.013), close to what proportional reallocation of the lost small-market procedures implies,[16] which fits patient travel as well as procedures continuing locally under a distant billing address.

Panel B of Table 8 turns to hospital outcomes around switches to documented adopting platforms. The share of abdominal CTs scanned with and without contrast, a protocoling decision involving radiologists, does not change; the share of emergency head CTs read within 45 minutes falls, but the fall begins two to three years before the switch (Appendix Table C.2, panel C), and at a single site Schmuelling et al. (2021) likewise found no change in report or turnaround times after a detection tool went live; two ordering measures set by referring clinicians do not move. The measures cover limited dimensions of care, and none shows a change at the switch.

In sum, procedures requiring a radiologist's presence left exposed small markets with interpretation, about half of the decline being work no longer performed there; whether access to care deteriorated, the data cannot establish.

(+0.04, s.e. 0.06).

[16] Log procedures in markets with at least 25 radiologists on the procedure-weighted mean exposure of small markets within the band (centroids from the 2017 Census gazetteer, U.S. Census Bureau 2017), standardized and interacted with the clock: +0.018 (0.012) within 100 kilometers, +0.028 (0.013) within 150, +0.014 (0.013) within 200. Proportional reallocation of the lost small-market volume implies about 0.028.

7 Quantifying the Reallocation of Radiology Production

Sections 5 and 6 documented a geographic and organizational reallocation of radiology production. We now use a parsimonious quantitative implementation of the model to examine how this reallocation interacts with staffing indivisibilities and minimum practice viability. The exercise addresses three questions: why local billing declines most sharply in thin markets; how much reader capacity can be saved by pooling production across establishments; and how organizational reallocation differs from an equivalent improvement in task productivity when production locations remain fixed.

The calibration is disciplined by observed geographic affiliations and billing patterns, rather than direct measures of AI verification. It therefore quantifies the implications of organizational reallocation under specified production and staffing assumptions, without structurally estimating the verification technology, agency rents, or the causal contribution of AI. Its value lies in connecting changes in organizational boundaries to resource requirements and local practice viability, which cannot be inferred directly from the reduced-form estimates. Appendix D describes the construction, calibration, and sensitivity analysis.

7.1 Accounting Framework and Model Fit

We take 2013 as the base year and 2024 as the end of the transition. Initial production in market m, V_m^0, consists of interpretations billed from the market, including facility and office settings. Annual reader capacity, $\bar{h}$, is set to 3,984 interpretations, the 75th percentile of the 2013 distribution. We allocate production equally across hospitals in each market, with each hospital requiring $n(v) = \max\{1, \lceil v/\bar{h} \rceil\}$ readers. The resulting market-level excess-capacity measure is

$$\pi_m = \frac{\bar{h} \sum_j n(V_{jm}^0)}{V_m^0} - 1.$$

We represent the propensity for production to cross market boundaries by

$$q_m^t = \Lambda(\alpha_t + \lambda \pi_m),$$

where Λ is the logistic function. The 2017 cross-section of hospital affiliations, measured by the share of affiliated radiologists located outside the hospital's CBSA, disciplines λ and α_0. The increase in this share from 0.24 in 2017 to 0.34 in 2025 disciplines α_1. We map the resulting change in organizational propensity into locally retained production through

$$r_m = \exp\{-\psi(q_m^1 - q_m^0)\}.$$

Affiliations discipline the geographic propensity, while billing changes discipline the mapping into production reallocation.

Interpretations no longer retained locally are allocated to markets with at least 25 radiologists,

in proportion to their 2013 radiologist headcounts. This captures the geographic concentration of billing rather than the physical location of interpretation. We also allow local practices to exit when retained production per hospital cannot support one reader at the administered reimbursement rate. The viability threshold is set at one reader-load, corresponding to a fee that covers the cost of a fully utilized reader; it is not fitted. Remaining interpretations from exiting practices are reallocated to larger markets.

Demand evolves with imaging events per Medicare fee-for-service beneficiary and fee-for-service enrollment. Nationally, events per beneficiary increase by approximately 9 percent while enrollment falls by approximately 10 percent, leaving aggregate measured demand roughly unchanged. We choose ψ to match the decline in interpretations billed from markets with fewer than 25 radiologists, to 0.63 of their 2013 level. The separate responses of the two smallest market-size classes, larger-market outcomes, and aggregate production are not targeted.

Table 9 reports the fit. Although only the combined decline for markets with fewer than 25 radiologists is targeted, the model reproduces the gradient within this group: retained production is 0.55 for markets with fewer than ten radiologists, versus 0.53 in the data, and 0.68 for markets with 10–24 radiologists, versus 0.69. The model also reallocates production toward larger markets: markets with at least fifty radiologists reach 1.10 of initial production, versus 1.09 in the data, and aggregate production is 0.99 of its initial level, against 0.98; the one class the model overpredicts is 25–49 radiologists (0.92 against 0.87).[17]

Volume-weighted excess capacity is 0.21 in markets with fewer than ten radiologists, compared with 0.07–0.08 elsewhere (Appendix Table D.2). Yet cross-market affiliations were already most prevalent in the smallest markets: in 2017, 70 percent of radiologists affiliated with their hospitals were located outside the CBSA, compared with 15 percent in the largest markets (Table 9, panel D). This share subsequently rose by only ten percentage points in the smallest class, versus 19 and 16 points in the middle classes. The sharper decline in small-market billing cannot be attributed simply to a larger increase in measured cross-market affiliations.

The minimum viability threshold amplifies reallocation in thin markets. Without practice exit, the model predicts retained production of 0.60 and 0.65 in the two smallest market classes, compared with 0.55 and 0.68 when exit is allowed.[18] In 2013, median production per hospital in markets with fewer than ten radiologists was approximately two reader-loads, and 30 percent of hospitals operated below one reader-load; the corresponding median in larger markets was five to eight reader-loads (Appendix Table D.1). Reducing locally retained volume can push thin production units below the viability threshold without similarly affecting larger units. In the calibrated model, exiting practices account for 21 percent of initial local reader positions in markets with fewer than ten radiologists, 2 percent in markets with 10–24 radiologists, 1 percent in markets with 25–49 radiologists, and none

[17] Appendix D.5 examines alternative rules for assigning displaced interpretations to hub markets and finds similar geographic reallocation patterns.

[18] The comparison recalibrates ψ in each specification to match the same aggregate retention target for markets with fewer than 25 radiologists. Because exit removes proportionately more volume from the smallest markets, allowing exit requires a lower ψ (3.10 rather than 3.78). The weaker pooling response then increases retention in the 10–24 class, where exit matters less. Holding ψ fixed, introducing exit weakly reduces retained production in every market.

Table 9: Calibration and model fit

	Data	Model
A. Targeted moments		
Cross-market affiliation share, 2017	0.24	0.24
Cross-market affiliation share, 2025	0.34	0.34
Reads billed from markets under 25 radiologists, 2024/2013	0.63	0.63
B. Untargeted outcomes: reads billed from the market, 2024/2013, by 2013 market size		
Fewer than 10 radiologists	0.53	0.55
10–24	0.69	0.68
25–49	0.87	0.92
50 or more	1.09	1.10
All markets	0.98	0.99
C. Share of 2013 local reader positions in units below the viability threshold (model)		
Fewer than 10 radiologists	—	0.21
of which below it on demand alone	—	0.13
10–24	—	0.02
25–49	—	0.01
50 or more	—	0.00

D. Cross-market affiliation share by 2013 market size, percent (data)

2013 market size	2017	2025	Change (points)
Fewer than 10 radiologists	70.3	80.2	+9.9
10–24	40.4	59.6	+19.2
25–49	36.7	52.7	+16.0
50 or more	14.9	22.4	+7.5

NOTE.—$\bar{h}$: 75th percentile of 2013 interpretations per billing radiologist (3,984); minimum complement one. Cross-market affiliation: share of hospital-affiliated radiologists located outside the hospital's CBSA in the Physician Compare and Doctors & Clinicians files. The 2017 cross-section disciplines λ and α_0, the 2025 share α_1 and the under-25 ratio ψ (values in Appendix Table D.1); the two under-25 classes are targeted jointly, nothing else is. Reads: all imaging interpretations billed from the market; reorganized reads are billed from markets with 25 or more radiologists in proportion to 2013 headcount; demand grows with imaging events per fee-for-service beneficiary times beneficiaries. Panel C: model-implied positions in units whose retained volume falls below one reader-load, not observed exits. Panel D: data only; only the aggregate shares in panel A are targeted.
SOURCE.—Authors' calculations from the calibrated model of Section 7.

in the largest; 13 of the 21 points are in units below the threshold on demand alone, which would exit without any reorganization (Table 9, panel C).

The exercise thus separates two economically distinct margins. Staffing indivisibilities create potential capacity savings from pooling, while minimum practice viability amplifies the geographic consequences of reallocation. The latter helps account for the disproportionate decline in the smallest markets, a pattern not separately targeted by the calibration. It is also consistent with the individual adjustments documented in Subsection 6.2: approximately 30 percent of small-market radiologists cease Medicare billing, 11 percent relocate their billing, and those remaining locally reduce their billed interpretations. Because the model tracks production units and reader positions rather than individual careers, these are comparisons of economic magnitudes and geographic incidence, not fitted worker transitions.

7.2 Capacity Savings and Counterfactual Allocations

Table 10 reports three counterfactual exercises using the calibrated model. They quantify how production indivisibilities affect geographic reallocation, how pooling changes aggregate reader-capacity requirements, and how these requirements respond to demand growth. The exercises hold calibrated parameters fixed unless otherwise indicated.

Production indivisibilities. Removing both discrete staffing and the minimum practice-viability threshold increases retained production in markets with fewer than ten radiologists from 0.55 to 0.65. Removing staffing indivisibilities alone leaves it at 0.55 (Table 10). The viability threshold, rather than excess staffing capacity by itself, accounts for most of the additional geographic decline in the smallest markets.

Table 10: Counterfactual implications

A. Production indivisibilities	Benchmark	Divisible production
Reads billed from the market, 2024/2013: fewer than 10 radiologists	0.55	0.65
10–24	0.68	0.69
25–49	0.92	0.92
50 or more	1.10	1.09
Same, divisible staffing only ($\pi_m = 0$), exit margin kept		
fewer than 10 radiologists	0.55	0.55
10–24	0.68	0.69
25–49	0.92	0.92
50 or more	1.10	1.10
Excess reader capacity in 2013	8.0%	0.0%
Capacity saving from reorganization, 2024	2.5%	—
B. Verification versus task productivity, equal resource saving	Verification	Task productivity
Aggregate resource saving	2.5%	2.5%
Implied task-productivity gain	—	2.5%
Production reorganized across markets	26.7%	0.0%
Reads billed from the market, 2024/2013: fewer than 10 radiologists	0.55	0.99
C. Demand growth at which aggregate reader requirements are unchanged		Break-even growth
Verification only		2.55%
Task productivity only		2.55%
Both		5.16%
Verification only, practice viability recomputed at each growth rate		2.55%

NOTE.—Panel A sets every market's excess capacity to zero, parameters fixed; the divisible-staffing-only rows keep the one-reader viability threshold. Excess capacity is discrete-staffing capacity relative to divisible production; the saving is the reduction in reader requirements relative to serving 2024 demand within the original local units. Panel B: a task-productivity gain chosen to give the same aggregate saving at fixed demand, with no reorganization, so reads follow demand in every market. Panel C: common demand growth at which aggregate reader requirements return to their 2013 level; verification and task productivity coincide by construction; the last row recomputes practice viability at each growth rate.
SOURCE.—Authors' calculations from the calibrated model of Section 7.

Staffing indivisibilities nevertheless matter for resource use. In 2013, discrete local staffing required 8.0 percent more reader capacity than divisible production. By reallocating approximately 27 percent of production across markets, the calibrated reorganization reduces reader-capacity requirements for 2024 demand by 2.5 percent relative to production within the original local units.[19]

Equivalent capacity savings with fixed production locations. To isolate the geographic consequences of organizational reallocation, we compare it with a task-productivity improvement that leaves production in its original markets. An increase in reading productivity of approximately 2.5 percent generates the same 2.5 percent reduction in aggregate reader-capacity requirements at fixed demand as the calibrated reorganization.

The two scenarios deliver equal aggregate savings by construction, but different geographic allocations. Under reorganization, locally billed interpretations in markets with fewer than ten radiologists fall to 0.55 of their 2013 level. Under the fixed-location productivity benchmark, interpretations instead follow local demand, with retained production of approximately 0.99 across market-size classes. The comparison illustrates how the same aggregate capacity saving can accompany sharply different local outcomes. It does not rule out organizational reallocation in response to productivity improvements when production boundaries are allowed to adjust.

Demand growth. We finally calculate the proportional increase in demand needed to restore aggregate reader-capacity requirements to their 2013 level. The threshold is approximately 2.6 percent under reorganization and, by construction, the same under the equivalent task-productivity benchmark; it rises to 5.2 percent when both improvements are combined. These calculations are conditional on the model's production and viability rules, rather than general predictions about employment.

Measured demand was roughly flat over the period, so the model implies a modest reduction in capacity requirements for the activities it covers; the 16 percent growth in Medicare-billing radiologists is not comparable, since it spans all payers and activities.

8 Conclusion

The prediction that artificial intelligence would make radiologists redundant focused on machines substituting for professional expertise. The transformation documented in this paper has been different. As reader-facing AI diffused, radiology production became increasingly organized across geographic markets, with declining local billing in small markets and expanding multi-market platforms, but without corresponding increases in measured interpretations per radiologist. We develop a framework in which AI can facilitate this reorganization by improving the verification of

[19] Recomputing whole-reader staffing after reallocation reduces the estimated saving from 2.5 to 1.5 percent. Using facility-setting interpretations alone, with the parameters recalibrated, yields savings of 3.3 percent under the baseline accounting rule and 0.7 percent with recomputed staffing. Thus, capacity savings range from 0.7 to 3.3 percent across these specifications (Appendix D.5).

professional output and reducing the costs of contracting for work performed at a distance. When production requires indivisible local staffing, improved verification can shift work toward organizations that pool demand across establishments, even without replacing professionals or substantially increasing individual throughput.

The empirical evidence is consistent with this mechanism. Local radiologist billing declines most sharply in smaller markets, while multi-market organizations expand their geographic reach; pathologists, who serve the same hospitals without the technology, show no such gradient. The geographic shift is also evident in where interpretations are physically performed across Medicare payment localities, and hospitals switching to platform coverage increasingly rely on remotely located radiologists. We find little evidence of increased measured interpretation throughput, including in imaging modalities directly covered by AI. A calibrated model of production reallocation reproduces the differences in local retention across market sizes and implies that substantial geographic reorganization can coexist with modest aggregate reader-capacity savings. In the baseline calibration, approximately 27 percent of production is reallocated across markets, reducing aggregate capacity requirements by 2.5 percent. An equally capacity-saving task-productivity improvement, holding production locations fixed, leaves the geographic allocation of work largely unchanged by construction.

The consequences extend beyond image interpretation. Radiologists initially billing from small markets were more likely to cease Medicare billing or relocate their billing activity, while those remaining locally experienced declines in billed interpretations. Procedures requiring radiologist involvement also became less likely to be billed from small markets, without compensating increases in local billing by other providers. Yet important limitations remain. Our research design identifies differential changes associated with predetermined exposure during national AI diffusion, not the causal effect of adoption by individual practices. We cannot measure radiologists' earnings, establish at the hospital level where procedures were physically performed, or determine whether patients' access to care changed. Nor can we uniquely isolate improved verification from other forces facilitating remote production.

Labor-market effects of AI cannot always be inferred from the tasks machines can perform. In expert services, production also depends on how professional judgment is certified, how work is coordinated, and whether organizations can pool staffing capacity across establishments. The importance of this channel depends on institutional arrangements governing remote delivery, professional certification, and the credibility of information about work performed elsewhere. As these arrangements evolve, other margins of adjustment, including direct task substitution, may become feasible. The economic consequences of AI depend not only on what machines can do, but also on how their capabilities change the organizations through which human expertise is produced, verified, and delivered.

References

Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy 40*(121), 13–58.

Acemoglu, D. and P. Restrepo (2019a). Artificial intelligence, automation, and work. In A. Agrawal, J. Gans, and A. Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda*, pp. 197–236. Chicago: University of Chicago Press.

Acemoglu, D. and P. Restrepo (2019b). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives 33*(2), 3–30.

Agarwal, N., A. Moehring, P. Rajpurkar, and T. Salz (2023). Combining human expertise with artificial intelligence: Experimental evidence from radiology. Working Paper 31422, National Bureau of Economic Research.

Agha, L., K. M. Ericson, and X. Zhao (2023). The impact of organizational boundaries on health care coordination and utilization. *American Economic Journal: Economic Policy 15*(3), 184–214.

Allen, B., S. Agarwal, L. Coombs, C. Wald, and K. Dreyer (2021). 2020 ACR Data Science Institute artificial intelligence survey. *Journal of the American College of Radiology 18*(8), 1153–1159.

Antràs, P., L. Garicano, and E. Rossi-Hansberg (2006). Offshoring in a knowledge economy. *Quarterly Journal of Economics 121*(1), 31–77.

Arrow, K. J. (1963). Uncertainty and the welfare economics of medical care. *American Economic Review 53*(5), 941–973.

Autor, D. H. (2015). Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives 29*(3), 3–30.

Baker, G., R. Gibbons, and K. J. Murphy (2002). Relational contracts and the theory of the firm. *Quarterly Journal of Economics 117*(1), 39–84.

Baker, G. P. and T. N. Hubbard (2003). Make versus buy in trucking: Asset ownership, job design, and information. *American Economic Review 93*(3), 551–572.

Baker, G. P. and T. N. Hubbard (2004). Contractibility and asset ownership: On-board computers and governance in U.S. trucking. *Quarterly Journal of Economics 119*(4), 1443–1479.

Bloom, N., L. Garicano, R. Sadun, and J. Van Reenen (2014). The distinct effects of information technology and communication technology on firm organization. *Management Science 60*(12), 2859–2885.

Brynjolfsson, E., D. Li, and L. Raymond (2025). Generative AI at work. *Quarterly Journal of Economics 140*(2), 889–942.

Capps, C., D. Dranove, and C. Ody (2018). The effect of hospital acquisitions of physician practices on prices and spending. *Journal of Health Economics 59*, 139–152.

Carmona, G. and K. Laohakunakorn (2024). Improving the organization of knowledge in production by screening problems. *Journal of Political Economy 132*(4), 1290–1326.

Chen, M., Y. Wang, Q. Wang, J. Shi, H. Wang, Z. Ye, P. Xue, and Y. Qiao (2024). Impact of human and artificial intelligence collaboration on workload reduction in medical image interpretation. *npj Digital Medicine 7*(1), 349.

Christensen, E. W., Y. Chung, E. Y. Rula, and J. R. Parikh (2024). Changes in the radiology practice landscape and indicators of practice consolidation from 2014 to 2023. *American Journal of Roentgenology 223*(2), e2431357.

Coase, R. H. (1937). The nature of the firm. *Economica 4*(16), 386–405.

Darby, M. R. and E. Karni (1973). Free competition and the optimal amount of fraud. *Journal of Law and Economics 16*(1), 67–88.

De Ridder, M. (2024). Market power and innovation in the intangible economy. *American Economic Review 114*(1), 199–251.

Dingel, J. I., J. D. Gottlieb, M. Lozinski, and P. Mourot (2023). Market size and trade in medical services. Working Paper 31030, National Bureau of Economic Research.

Dogra, S., X. Zhang, I. Silva, Ezequiel, and P. Rajpurkar (2025). The financial, operational, and clinical advantages of generalist radiology AI. *Radiology 316*(3), e242362.

Dranove, D. and C. Ody (2019). Employed for higher pay? how Medicare payment rules affect hospital employment of physicians. *American Economic Journal: Economic Policy 11*(4), 249–271.

Dulleck, U. and R. Kerschbamer (2006). On doctors, mechanics, and computer specialists: The economics of credence goods. *Journal of Economic Literature 44*(1), 5–42.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science 384*(6702), 1306–1308.

Frey, C. B. and M. A. Osborne (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change 114*, 254–280.

Garicano, L. (2000). Hierarchies and the organization of knowledge in production. *Journal of Political Economy 108*(5), 874–904.

Garicano, L. and T. N. Hubbard (2007). Managerial leverage is limited by the extent of the market: Hierarchies, specialization, and the utilization of lawyers' human capital. *Journal of Law and Economics 50*(1), 1–43.

Garicano, L. and E. Rossi-Hansberg (2006). Organization and inequality in a knowledge economy. *Quarterly Journal of Economics 121*(4), 1383–1435.

Garicano, L. and E. Rossi-Hansberg (2015). Knowledge-based hierarchies: Using organizations to understand the economy. *Annual Review of Economics 7*(1), 1–30.

Geis, J. R., A. P. Brady, C. C. Wu, J. Spencer, E. Ranschaert, J. L. Jaremko, S. G. Langer, A. Borondy Kitts, J. Birch, W. F. Shields, R. van den Hoven van Genderen, E. Kotter, J. Wawira Gichoya, T. S. Cook, M. B. Morgan, A. Tang, N. M. Safdar, and M. Kohli (2019). Ethics of artificial intelligence in radiology: Summary of the joint European and North American multisociety statement. *Radiology 293*(2), 436–440.

Goldfarb, A. and C. Tucker (2019). Digital economics. *Journal of Economic Literature 57*(1), 3–43.

Goldsmith-Pinkham, P., C. Tan, and A. K. Zentefis (2026, January). Human-AI collaboration in radiology: The case of pulmonary embolism. Working Paper 2601.13379, arXiv.

Grossman, S. J. and O. D. Hart (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy 94*(4), 691–719.

Gu, Z., S. Dogra, M. Siriruchatanon, J. Kneifati-Hayek, and S. K. Kang (2026). Radiology workflow assistance with artificial intelligence: Establishing the link to outcomes. *Journal of the American College of Radiology 23*(3), 389–398.

Hampole, M., D. Papanikolaou, L. D. W. Schmidt, and B. Seegmiller (2025). Artificial intelligence and the labor market. Working Paper 33509, National Bureau of Economic Research.

Hanna, T. N., S. D. Steenburg, A. B. Rosenkrantz, R. S. Pyatt, R. Duszak, and E. B. Friedberg (2020). Emerging challenges and opportunities in the evolution of teleradiology. *American Journal of Roentgenology 215*(6), 1411–1416.

Hsieh, C.-T. and E. Rossi-Hansberg (2023). The industrial revolution in services. *Journal of Political Economy Macroeconomics 1*(1), 3–42.

Humlum, A. and E. Vestergaard (2025). Still waters, rapid currents: Early labor market transformation under generative AI. Working Paper 33777, National Bureau of Economic Research.

Ide, E. and E. Talamàs (2025). Artificial intelligence in the knowledge economy. *Journal of Political Economy 133*(12), 3762–3800.

Innes, R. D. (1990). Limited liability and incentive contracting with ex-ante action choices. *Journal of Economic Theory 52*(1), 45–67.

Khunte, M. and Y. Singh (2025). Private equity acquisitions of radiology practices from 2013 to 2023: National- and state-level analyses. *American Journal of Roentgenology 224*(5), e2532738.

Klein, B. and K. B. Leffler (1981). The role of market forces in assuring contractual performance. *Journal of Political Economy 89*(4), 615–641.

Lång, K., V. Josefsson, A.-M. Larsson, S. Larsson, C. Högberg, H. Sartor, S. Hofvind, I. Andersson, and A. Rosso (2023). Artificial intelligence-supported screen reading versus standard double reading in the mammography screening with artificial intelligence trial (MASAI): A clinical safety analysis of a randomised, controlled, non-inferiority, single-blinded, screening accuracy study. *The Lancet Oncology 24*(8), 936–944.

Langlotz, C. P. (2019). Will artificial intelligence replace radiologists? *Radiology: Artificial Intelligence 1*(3), e190058.

Lashkari, D., A. Bauer, and J. Boussard (2024). Information technology and returns to scale. *American Economic Review 114*(6), 1769–1815.

Levin, J. (2003). Relational incentive contracts. *American Economic Review 93*(3), 835–857.

MacLeod, W. B. (2003). Optimal contracting with subjective evaluation. *American Economic Review 93*(1), 216–240.

McElheran, K., J. F. Li, E. Brynjolfsson, Z. Kroff, E. Dinlersoz, L. Foster, and N. Zolas (2024). AI adoption in america: Who, what, and where. *Journal of Economics & Management Strategy 33*(2), 375–415.

Pettet, G., J. West, D. Robert, A. Khetani, S. Kumar, S. Golla, and R. Lavis (2024). A retrospective audit of an artificial intelligence software for the detection of intracranial haemorrhage used by a teleradiology company in the United Kingdom. *BJR—Open 6*(1), tzae033.

Rosenkrantz, A. B., H. B. Fleishon, I. Silva, Ezequiel, C. E. Bender, and J. Duszak, Richard (2020). Radiology practice consolidation: Fewer but bigger groups over time. *Journal of the American College of Radiology 17*(3), 340–348.

Schmuelling, L., F. C. Franzeck, C. H. Nickel, G. Mansella, R. Bingisser, N. Schmidt, B. Stieltjes, J. Bremerich, A. W. Sauter, T. Weikert, and G. Sommer (2021). Deep learning-based automated detection of pulmonary embolism on CT pulmonary angiograms: No significant effects on report communication times and patient turnaround in the emergency department nine months after technical implementation. *European Journal of Radiology 141*, 109816.

Silva, Ezequiel, I., J. Breslau, R. M. Barr, L. A. Liebscher, M. Bohl, T. Hoffman, G. W. L. Boland, C. Sherry, W. Kim, S. S. Shah, and M. Tilkin (2013). ACR white paper on teleradiology practice: A report from the task force on teleradiology practice. *Journal of the American College of Radiology 10*(8), 575–585.

Steinbrook, R. (2007). The age of teleradiology. *New England Journal of Medicine 357*(1), 5–7.

Topff, L., S. Steltenpool, E. R. Ranschaert, N. Ramanauskas, R. Menezes, J. J. Visser, R. G. H. Beets-Tan, and N. S. Hartkamp (2024). Artificial intelligence-assisted double reading of chest radiographs to detect clinically relevant missed findings: A two-centre evaluation. *European Radiology 34*(9), 5876–5885.

van Leeuwen, K. G., S. Schalekamp, M. J. C. M. Rutten, B. van Ginneken, and M. de Rooij (2021). Artificial intelligence in radiology: 100 commercially available products and their scientific evidence. *European Radiology 31*(6), 3797–3804.

Webb, M. (2020). The impact of artificial intelligence on the labor market. Working paper, Stanford University.

Wenderott, K., J. Krups, F. Zaruchas, and M. Weigl (2024). Effects of artificial intelligence implementation on efficiency in medical imaging—a systematic literature review and meta-analysis. *npj Digital Medicine 7*(1), 265.

Yu, F., A. Moehring, O. Banerjee, T. Salz, N. Agarwal, and P. Rajpurkar (2024). Heterogeneity and predictors of the effects of AI assistance on radiologists. *Nature Medicine 30*(3), 837–849.

Zhang, K., B. Khosravi, S. Vahdati, and B. J. Erickson (2024). FDA review of radiologic AI algorithms: Process and challenges. *Radiology 310*(1), e230242.

Online Appendix

AI and the Boundary of the Firm: Evidence from Radiology

Marco Francesconi, Xiaoyu Xia

A. Supplementary Information on the Data i

B. Exposure: Construction, Mean Reversion and Robustness iv

C. Additional Results ix

D. Quantitative Model and Counterfactuals xvi

E. Data and Documentary Sources xxiii

A Supplementary Information on the Data

A.1 Sources and Where Used

Table A.1 lists every source. The first five are described in Section 4; the others are named where they are first used.

Table A.1: Data sources and where each is used

Source	Measures	Used in
FDA AI-enabled device list, radiology panel, to March 2026 (2026)	national and modality diffusion clocks (reader-facing devices)	§4.1
Medicare Physician & Other Practitioners files, 2013–24 (2026g); 2012 from the NBER mirror (2026)	radiologists billing; interpretations, services, payments; procedures; comparison specialties	§4.2; §5–6
Physician/Supplier Procedure Summary, 2016–24 (2026i), with the ZIP–locality files (2024)	interpretations and procedures by payment locality (where read or performed)	§4.2; Tables 3, C.4
Physician Compare / Doctors & Clinicians, five snapshots 2017–25 (2025b)	groups (mapped to parent organizations), practice locations, hospital affiliations, graduation year	§4.3; §5–6
Medicare Geographic Variation files, 2007–24 (2026d)	imaging received, by beneficiary residence	§4.2; §5.3
AI adoption record (vendor and firm announcements), 50 events, 24 organizations, 2018–26	who adopted, what and when	§4.3; §5.4
Census population estimates (2025), OMB delineation (2018) and gazetteer (2017); HUD ZIP–CBSA crosswalk (2017)	markets, size control, distances	§4.2; §6.3
BLS Occupational Employment and Wage Statistics (2026)	imaging technologists, metropolitan areas	§5.6
CMS Hospital General Information, 2019–26 (2026b)	hospitals by market; closures	Table A.3; §5.6
CMS Hospital Compare (2026a)	head-CT turnaround (OP-23); protocoling and ordering (OP-10, OP-8, OP-13)	§5.5; §6.3
Medicare Inpatient file (2026e)	thrombectomy volumes (DRG 023, 024)	§5.5
Medicare Advantage enrollment by county (2026f)	Medicare Advantage penetration	Table B.2
Interstate Medical Licensure Compact Commission (2024)	states' operational dates	§5.6

NOTE.—HCPCS codes in the hospital-setting exposure basket: 71020, 72100, 70450, 71250, 74176, 74150, 70551, 72148, 76700, 76770, 76856, 77057, G0202, 78306, 78815. Interpretations: HCPCS 70010–79999 and the mammography G-codes, excluding the computer-aided-detection add-on codes 77051 and 77052, billed beside mammogram reads until 2016 and bundled into the mammography codes from 2017. Pathology exposure basket (specimen codes): 88304, 88305, 88307, 88312, 88341, 88342. Geographic Variation: the historical (BETOS, 2007–13) and current (RBCS, 2014–24) series align across the 2014 change: county fee-for-service beneficiary counts correlate 0.997 and county imaging rates 0.946, against 0.962 between adjacent years within the current series. Procedure basket (Section 6.3): 10005, 10021, 10022, 19081, 19083, 32555, 36569, 36573, 47000, 49083, 49406, 62270, 62328, 77012; fluoroscopic studies 74220, 74230, 74240, 74270.

A.2 Device Classification

Devices on the FDA's radiology panel are classified by product code and device description into the stage of the imaging pipeline they act on; reader-facing devices analyze the content of an image and report on it. Table A.2 gives the counts by pipeline stage, function and task; the cumulative series is in Figure 2.

Table A.2: FDA-cleared radiology AI devices by pipeline stage, function and task

Function	Devices
Reader-facing (analyses image content)	249
Reader-facing, dental imaging (set aside)	30
Image processing and quantification	406
Scanner-embedded acquisition and reconstruction	399
Radiotherapy planning	79
Other	1
Total	1,164
Reader-facing devices by function (FDA product-code classes, Appendix A.2)	
notification and triage	118
second-reader detection and characterization	99
quantification	32
Reader-facing devices by task	
ICH triage	22
breast cancer detection (mammo)	17
LVO triage	14
breast cancer detection (DBT)	11
pneumothorax triage	11
all other tasks	174

NOTE.—Devices are classified by product code and device description. Classes name the stage of the imaging pipeline a device acts on, not a legal status. Reader-facing devices analyse the content of an image; none is cleared to issue a report without a licensed reader's sign-off. The dental devices, which radiologists do not use, are set aside from the reader-facing class, the clocks and the figures. ICH: intracranial hemorrhage; LVO: large-vessel occlusion; DBT: digital breast tomosynthesis; mammo: mammography.
SOURCE.—Authors' classification of the FDA list of AI-enabled medical devices, radiology panel, through March 2026.

A.3 Denominators by Market Size

Table A.3 counts hospitals, critical-access hospitals, population and interpretation volume by the 2013 size class of the market, including the two classes outside the estimation sample; Subsection 4.4 draws on it.

Table A.3: Where the reorganization lands: hospitals, population, volume and radiologists by 2013 market size

		Count (percent of U.S.)			Percent of U.S.	
Market class	Markets	Hospitals	Critical access	Population (m)	Volume	Radiologists
Fewer than 10 radiologists	431	769 (17)	260 (19)	28.1 (9)	9	5
10–24	146	456 (10)	117 (8)	25.6 (8)	11	7
25 or more	187	2,375 (51)	222 (16)	243.8 (75)	80	88
No billing radiologist in 2013	185	231 (5)	113 (8)	9.3 (3)	—	—
Outside any CBSA	—	822 (18)	670 (48)	18.3 (6)	—	—

NOTE.—Markets: every CBSA with a billing radiologist in 2013 (764; 23 of them lack the population estimates the regressions need; the 769 markets of Table 2 are the other 741 plus 38 whose first radiologist bills after 2013, less ten observed in a single year); the 25-or-more row pools the two largest classes; the last two rows are outside the estimation sample. Hospitals: acute-care and critical-access hospitals in the 2019 Hospital General Information file, assigned to markets by ZIP code. Population: 2017 estimate; the last row is the residual against the national total. Volume: 2013–14 hospital-setting interpretations, the volume that builds exposure. Radiologists: 2013 billing radiologists. Implied 2024 shortfall in above-average-exposure markets relative to average exposure: 842 billing radiologists (95 percent confidence interval 576–1,120), 439 of them in markets with fewer than 25 radiologists.

SOURCES.—Authors' calculations from the CMS Hospital General Information file; Census population estimates; the Medicare Physician & Other Practitioners files.

B Exposure: Construction, Mean Reversion and Robustness

B.1 Regression to the Mean: Simulation and Decomposition

Because exposure is built from billing volume, and billing volume is mechanically related to the number of radiologists billing, a market with transitorily many radiologists in the exposure window will have high measured exposure and will subsequently fall back. This is the most serious threat to the design, and we address it with a calibrated simulation rather than with placebos, since a mean-reversion artifact predicts the usual placebos as well.

We simulate markets with no technological change, matching the estimation sample exactly—the same 8,375 market-years, the same headcount distribution, the observed idiosyncratic variance (standard deviation 0.16 at the median market-year, seven radiologists) and the observed persistence (autocorrelation ≈ 0.80)—push each replication through the identical exposure-construction and estimation pipeline, and read off the coefficient. Across 5,000 replications, mean reversion generates between five and thirteen percent of the observed coefficient (three to six percent without the size control), and about two-fifths of it under an implausible worst case in which all headcount variation is Poisson noise. In the calibrated scenarios no replication produced a coefficient as negative as the observed one; in the Poisson worst case 58 of 5,000 did. The simulated artifact also has the wrong shape: it produces a hump *inside* the exposure window and then settles at a small negative constant, about -0.03 by 2024, rather than a decline that continues for a decade.

Three model-free checks agree. Decomposing exposure into persistent and transitory components loads the effect entirely on the persistent part (-0.059, $t = -6.2$) and not at all on the transitory part ($+0.013$, $t = 1.3$), the reverse of what mean reversion implies (Table B.2, panel A). Controlling directly for a market's headcount deviation during the exposure window—its 2013–14 log headcount less the 2012 value, interacted with the clock—leaves the coefficient at -0.080, with the deviation term at -0.002 (0.009). And building exposure from the hold-out window, 2013–14, rather than from 2015–17 gives a *larger* estimate, -0.088 ($t = -6.5$) against -0.052 ($t = -3.5$) with the size control, not a smaller one. We adopt the hold-out window as the paper's baseline because it does not overlap the outcome period.

B.2 Which Pre-Period Window Builds Exposure

Table B.1 rebuilds the exposure variable from every one-, two- and three-year window between 2013 and 2017 and re-estimates Equation (9) on each. Every window yields a negative gradient, and without the size control every estimate is at least five standard errors from zero. The variation across windows is systematic rather than noisy, and its direction is informative. Regression to the mean contaminates a window in proportion to its overlap with the outcome period, so it predicts that windows ending later—which share more years with the outcome deviation—should carry *larger* spurious declines. They carry smaller ones: the coefficient attenuates as the window's end year moves later, from -0.123 on 2013–15 to -0.069 on 2017 without the size control, and from -0.097 to -0.023 with it. The correlation between window end-year and coefficient is $+0.66$ and

+0.76 respectively, where the mean-reversion account requires a negative one.

We read the attenuation as measurement contamination of the opposite kind. Platform organization was diffusing before the 2018–19 clearance wave—platform formation is already visible in the 2017 affiliation data—so volume measured in 2016 or 2017 is measured after the reallocation has begun, and exposure built from it is a partly post-treatment variable. This is the reason to build exposure early, and it is also the honest limit of the exercise: the size-controlled estimate on the latest windows falls to -0.04 ($t = -2.6$) for 2016–17 and -0.02 ($t = -1.7$), no longer significant, for 2017, so the result does require measuring exposure before the reorganization is underway.

Table B.1: The exposure gradient across pre-period windows

	No size control		Size-controlled	
Window building exposure	coef.	(s.e.)	coef.	(s.e.)
2013	−0.107	(0.012)	−0.080	(0.015)
2013–14 †	−0.115	(0.011)	−0.088	(0.013)
2013–15	−0.123	(0.011)	−0.097	(0.016)
2014	−0.103	(0.010)	−0.076	(0.012)
2014–15	−0.109	(0.011)	−0.082	(0.013)
2014–16	−0.108	(0.011)	−0.076	(0.014)
2015	−0.100	(0.010)	−0.071	(0.013)
2015–16	−0.100	(0.011)	−0.066	(0.014)
2015–17	−0.095	(0.011)	−0.052	(0.015)
2016	−0.088	(0.011)	−0.051	(0.014)
2016–17	−0.084	(0.012)	−0.037	(0.014)
2017	−0.069	(0.011)	−0.023	(0.013)

NOTE.—Each row rebuilds exposure from hospital-setting interpretation volume over the stated window and re-estimates Equation (9). Dependent variable: log radiologists billing from the market, 2013–2024; $N = 8{,}375$ market-years in 769 markets for every row; the size control is ln population × clock. † marks the window used in the main text.
Standard errors in parentheses, clustered by market.
SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files; Census population estimates.

B.3 Robustness of the Gradient

Table B.2 collects the specification battery on the baseline construction: Medicare Advantage penetration and pre-period radiologist density interacted with the clock, the pandemic years dropped, alternative forms of the reader-facing clock (-0.088 to -0.091) and a clock of every radiology-panel AI device (-0.087), count and outcome variants, and randomization inference, under which none of 499 permutations of exposure across markets produces a coefficient as large in absolute value as the observed one (the largest is 0.030).

Table B.2: Robustness of the exposure gradient

	coef.	(s.e.)	N
A. Specification			
Baseline (Table 2)	−0.088	(0.013)	8,375
+ Medicare Advantage penetration	−0.085	(0.014)	8,375
+ clock × pre-period MA penetration	−0.089	(0.014)	8,375
+ clock × pre-period (2012) radiologist density	−0.087	(0.018)	8,124
ln pop × year indicators in place of ln pop × clock	−0.087	(0.013)	8,375
dropping 2020	−0.089	(0.014)	7,697
dropping 2020–21	−0.089	(0.014)	7,022
dropping 2020–22	−0.089	(0.014)	6,353
baseline on the sample with graduation years	−0.086	(0.015)	8,152
+ clock × market mean graduation year	−0.092	(0.015)	8,152
+ clock × share graduated by 1985	−0.090	(0.015)	8,152
+ market PE share and exposure × PE share: exposure × clock	−0.087	(0.013)	8,375
exposure × PE share, same regression	−0.024	(0.095)	8,375
+ exposure × national PE share: exposure × clock	−0.095	(0.023)	8,375
2015–17 exposure split: persistent component (predicted by 2013–14)	−0.059	(0.010)	8,375
transitory component, same regression	0.013	(0.010)	8,375
+ clock × (2013–14 minus 2012 log headcount): exposure	−0.080	(0.017)	8,078
the deviation term	−0.002	(0.009)	8,078
B. Diffusion clock variants (baseline: log of cumulative reader-facing clearances)			
level of the count	−0.090	(0.014)	8,375
square root of the count	−0.091	(0.014)	8,375
log, mid-year count	−0.089	(0.014)	8,375
log, dental devices included	−0.088	(0.014)	8,375
all radiology-panel AI devices, log	−0.087	(0.013)	8,375
C. Estimator and outcome variants			
PPML on the headcount level (positive cells)	−0.048	(0.052)	8,375
Weighted by 2013 headcount	−0.015	(0.044)	8,184
Metropolitan areas only	−0.072	(0.032)	4,509
Micropolitan areas only	−0.093	(0.016)	3,866
PPML incl. zero cells	−0.055	(0.049)	9,218
small-market interaction	−0.100	(0.060)	9,218
$\ln(1+n)$ incl. zero cells	−0.082	(0.011)	9,218
LPM: P(any billing radiologist)	−0.046	(0.011)	9,218
Headcount by current billing address (App. C.3 construction)	−0.088	(0.013)	8,375
by frozen first-billing address (exit and entry only)	−0.029	(0.008)	8,414
ln payment per service	−0.004	(0.017)	8,375
ln interpretations per radiologist	−0.150	(0.038)	8,355
ln total payments billed from market	−0.169	(0.037)	8,375

NOTE.—Dependent variable, unless the row says otherwise: log radiologists billing from the market, 2013–2024; exposure is the 2013–14 measure and the size control ln population × clock is included. The 2015–17 split rows decompose 2015–17 exposure into the part predicted by 2013–14 exposure and the residual; the 2012-deviation rows add the clock interacted with the market's 2013–14 minus 2012 log headcount. PE share: share of a market's listed radiologists in private-equity-backed groups, by directory snapshot, from the ownership record; national PE share: the same share nationally by year; in markets with fewer than 10 radiologists the gradient with the PE terms is −0.123 (0.019). The PPML rows weight markets by their counts. The clock counts reader-facing devices (Appendix Table A.2), excluding 30 dental devices; the last row of panel B uses every radiology-panel AI device instead. PPML: Poisson pseudo-maximum likelihood; LPM: linear probability model; MA: Medicare Advantage.
Standard errors in parentheses, clustered by market.
SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files; Census population estimates; the FDA list of AI-enabled medical devices; Medicare Monthly Enrollment; the authors' record of group ownership; the Physician Compare and Doctors & Clinicians files.

Table B.3: Robustness of the exposure gradient, continued: by market class, and function-specific clocks

	coef.	(s.e.)	N
A. By 2013 market class: exposure × clock, and its increment after 2021			
Fewer than 10 radiologists	−0.129	(0.029)	4,266
× post-2021 increment	0.013	(0.046)	
10–24	−0.049	(0.103)	1,710
× post-2021 increment	0.153	(0.086)	
25 or more	0.083	(0.049)	2,208
× post-2021 increment	−0.053	(0.084)	
B. Function-specific diffusion clocks (log cumulative clearances of devices with that function)			
Second-reader devices, clock alone	−0.088	(0.014)	8,375
Triage devices, clock alone	−0.085	(0.013)	8,375
Quantification devices, clock alone	−0.076	(0.012)	8,375
Second-reader, three clocks jointly	−0.045	(0.040)	8,375
Triage, three clocks jointly	−0.040	(0.028)	8,375
Quantification, three clocks jointly	−0.003	(0.017)	8,375

NOTE.—Panel A: Equation (9) with the size control, estimated within each 2013 market class, with a second term for the clock after 2021; the 25-or-more row pools the two largest classes. Panel B: each function clock is the log of one plus the cumulative count of reader-facing devices with that function (Appendix Table A.2), interacted with exposure; the three clocks correlate 0.84 to 0.97 across years, so the joint estimates are not separately identified. Same dependent variable and sample as Appendix Table B.2.
Standard errors in parentheses, clustered by market.
SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files; the FDA list of AI-enabled medical devices; Census population estimates.

Table B.4: What the gradient loads on: the exposure measure against alternative constructions

	Construction		Baseline beside it		
Exposure construction (each × clock, with ln pop × clock)	coef.	(s.e.)	coef.	(s.e.)	N
A. The volume measure in other forms					
Hospital-setting volume, log, residualized on ln pop (baseline)	−0.088	(0.013)	—	—	8,375
raw log volume, not residualized	−0.112	(0.017)	—	—	8,375
log volume per capita	−0.094	(0.014)	—	—	8,375
B. Scale and workload alternatives, entered beside the baseline					
Mean log volume per billing site	0.033	(0.027)	−0.148	(0.057)	8,181
Largest billing site, log volume	0.051	(0.037)	−0.178	(0.071)	8,181
Log number of billing sites (beside raw log volume)	−0.016	(0.016)	−0.106	(0.029)	8,181
C. Demand-side constructions, from the beneficiary's residence, entered alone					
Imaging events per 1,000 beneficiaries, 2010–12	−0.022	(0.008)	—	—	8,375
Beneficiary imaging volume, log, residualized on ln pop	−0.034	(0.007)	—	—	8,375

NOTE.—Dependent variable: log radiologists billing from the market, 2013–2024; market and year fixed effects. Panel A re-estimates Equation (9) with volume in other forms; given the size control these are linear transformations of one another, so the three t-statistics coincide. Panel B enters each alternative beside the baseline measure (the number of sites beside raw log volume), both interacted with the clock. Panel C replaces exposure with a demand-side measure and keeps the size control; the imaging-events rate uses the 2010–12 historical Geographic Variation file, the closest pre-period window available in that series, and the beneficiary-volume row uses 2013–14.
Standard errors in parentheses, clustered by market.
SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files; the Medicare Geographic Variation files; Census population estimates.

Table B.5: The intensive and extensive margins, all count outcomes treated symmetrically

	Zero	Log (positive cells)			Balanced panel				
					$\ln(1+n)$		LPM: P(any)		
Outcome	cells	coef.	(s.e.)	N	coef.	(s.e.)	coef.	(s.e.)	N
Radiologists billing locally	9%	−0.088	(0.013)	8,375	−0.082	(0.011)	−0.046	(0.011)	9,218
Hospital interpretation volume billed locally	10%	−0.337	(0.050)	8,239	−0.564	(0.072)	−0.040	(0.012)	9,218
Control: pathologists (own exposure)	9%	0.009	(0.018)	6,613	0.038	(0.015)	0.053	(0.016)	7,283
markets with fewer than 10 pathologists	—	−0.004	(0.021)	4,585	—	—	—	—	
markets with 10 or more pathologists	—	0.054	(0.041)	2,028	—	—	—	—	
Placebo: family-practice physicians	0%	−0.004	(0.007)	9,209	−0.005	(0.007)	−0.001	(0.001)	9,218
Placebo: psychiatrists	15%	−0.012	(0.014)	7,826	−0.009	(0.009)	0.006	(0.006)	9,218
Imaging workforce, metropolitan employment (OES, 2013–2025; no zero cells, log specification only)									
Radiologic technologists	—	0.021	(0.008)	4,335	—	—	—	—	
Diagnostic medical sonographers	—	0.005	(0.015)	3,162	—	—	—	—	
Nuclear medicine technologists	—	−0.003	(0.012)	1,279	—	—	—	—	
All occupations	—	−0.011	(0.006)	4,910	—	—	—	—	

NOTE.—Each cell is the coefficient on exposure × clock from Equation (9) with the size control, on a balanced panel of estimation-sample markets, 2013–2024. "Zero cells" is the share of market-years in which the outcome is zero; the log specification drops them, the other two include them. Pathology uses its own exposure, built in the same way from 2013–14 hospital specimen volume. The imaging-workforce rows use Occupational Employment and Wage Statistics (OES) metropolitan employment, 2013–2025, which has no zero cells and covers the metropolitan areas with employment data (77 of the 431 markets with fewer than 10 radiologists). LPM: linear probability model.

Standard errors in parentheses, clustered by market.

SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files; BLS Occupational Employment and Wage Statistics; Census population estimates.

C Additional Results

C.1 Hospital Contract Flows

Hospital-level coverage is consistent with the make-or-buy move of Proposition 1. Among hospitals not already covered by an adopter-majority group in 2017, those in markets with fewer than ten radiologists switched at 13.6 percent, against 9.9 percent in markets with fifty or more, and switching is more likely at hospitals with fewer radiologists of their own (logit coefficient −0.17 per log point, $t = -2.9$; Table C.1). At the switching hospital, affiliated radiologists located in its own market fall by 0.29 log points and those located elsewhere rise by 0.29, leaving the total unchanged (panel B), though the local count was already drifting down before the switch. Market totals barely move: the count of single-market non-platform radiologists falls by 0.12 log points after a local switch (0.17 in markets with fewer than twenty-five radiologists), billing headcount by 0.003 ($t = -0.1$), and the event study is flat, because in larger markets most dropped readers stay and read for someone else. In markets with fewer than 25 radiologists, 52 percent of a switching hospital's local readers are no longer located there two years later, against 33 percent at hospitals that never switch.

Table C.1: Hospitals switching to adopter-majority coverage

	Estimate	(s.e.)	N
A. Which hospitals switch, 2017–2025			
Share switching: markets with fewer than 10 radiologists	0.136		605
10–49	0.101		671
50 or more	0.099		1,593
Logit on own radiologists, log 2017 [marginal effect]	−0.167 [−0.016]	(0.057)	2,911
B. The switching hospital: its affiliated radiologists after the switch			
Located in the hospital's market	−0.293	(0.044)	13,944
markets with fewer than 25 radiologists	−0.238	(0.068)	4,661
Located elsewhere	0.286	(0.064)	13,944
All affiliated radiologists	0.009	(0.045)	13,944
C. The market: radiologists after a local hospital switches			
Single-market non-platform radiologists	−0.120	(0.052)	3,882
markets with fewer than 25 radiologists	−0.174	(0.067)	2,492
Radiologists billing from the market	−0.003	(0.026)	5,464

NOTE.—A hospital switches when most of its affiliated radiologists become adopter-platform readers in a later directory snapshot; base: hospitals not adopter-majority in 2017. Panel A: shares by 2013 market size (N: hospitals); logit with market exposure and log population. Panel B: log of one plus the count, hospital and snapshot fixed effects; located means a primary practice location in the hospital's CBSA; the local count was already falling before the switch. Panel C: logs, market and year fixed effects; the first two rows use the five directory snapshots, the last the annual claims panel, 2013–2024; its event study is flat (all coefficients within ±0.02).

Standard errors in parentheses, clustered by market.

SOURCES.—Authors' calculations from the Physician Compare and Doctors & Clinicians files; the authors' record of documented AI adoption; the Medicare Physician & Other Practitioners files.

C.2 The Task-Margin Battery in Detail

Table 6 summarizes the task-margin battery; this appendix gives the rows that are not reported in full in the text.

Firms. Markets with a hospital covered by vRad, the teleradiology firm that had deployed AI by 2017 (vRad (Virtual Radiologic), 2020), are compared with markets covered by teleradiology firms with no AI (Table C.2, panel A). Neither coverage indicator carries a gradient once exposure and size are in the regression, and vRad coverage does not steepen the exposure gradient. In that period vRad supplied reads to local groups that kept the hospital contracts (Subsection 3.6): faster reading at a supplier that does not hold the contract leaves local positions in place.

The competitor margin. If the tools had made platform readers more productive, platforms could have taken work from the independent radiologists competing with them, whose numbers should then have fallen. Comparing the count of non-platform, single-market radiologists listed in markets where a documented adopter was present in 2017 or 2019 (102 and 220 markets) with other markets after 2021, with size-by-post and exposure-by-post controls, gives +0.21 (0.06); in markets with fewer than twenty-five radiologists, +0.17 to +0.19. The event study shows why this is not an effect: relative to 2021, the coefficients are −0.26 in 2017, −0.08 in 2019, +0.06 in 2023 and +0.06 in 2025, a rising trend in adopter-present markets that predates adoption and does not change at it.

Outcomes. Panel B of Table C.2 reports the head-CT turnaround comparison under four definitions of adopter coverage; none improves; the event study around 2019 is flat through 2024, with the March 2025 point at −3.3 (1.5). Panel C of Table C.2 turns to the hospitals whose coverage switched to an adopting platform. Turnaround there fell by about four points after the switch (−3.8, s.e. 1.6), but the event study shows the decline under way two to three years before it, so the estimate is better read as a trend the switch did not interrupt than as its effect. The protocoling measure does not move before or after the switch; the two ordering placebos do not move at the switch either, though the cardiac-imaging measure is 0.6 points lower three or more years before it (−0.64, 0.26).

The subsidy. From October 2020 to September 2022 Medicare paid a New Technology Add-on Payment of roughly $1,040 per use for AI detection of large-vessel occlusion in stroke (approved for fiscal 2021 and extended once; Centers for Medicare & Medicaid Services 2020b, 2021). If the payment changed what hospitals did, thrombectomy volumes at the hospitals it paid should rise when it switched on and fall when it switched off. Among the 332 hospitals performing thrombectomy in both 2018 and 2019, log discharges relative to 2019 (hospital fixed effects) sit between −0.08 and −0.12 in 2020–2022 and at −0.03 (0.03) in 2024: no rise during the window and no fall at its expiry. Head-CT turnaround at thrombectomy-capable hospitals relative to others

Table C.2: The task-margin battery in detail: teleradiology coverage, turnaround, and hospitals that switched

A. Markets covered by the AI-using teleradiology firm versus non-AI teleradiology firms

Term	coef.	(s.e.)	N
All markets			
Covered by vRad in 2017 × clock	0.012	(0.020)	8,375
Covered by a non-AI teleradiology firm in 2017 × clock	−0.026	(0.022)	8,375
Exposure × clock (same regression)	−0.088	(0.013)	8,375
vRad-covered × exposure × clock	0.010	(0.029)	8,375
Markets with fewer than 25 radiologists			
Covered by vRad in 2017 × clock	0.011	(0.035)	5,976
Covered by a non-AI teleradiology firm in 2017 × clock	−0.024	(0.039)	5,976
Exposure × clock (same regression)	−0.103	(0.016)	5,976
vRad-covered × exposure × clock	−0.032	(0.068)	5,976

B. Emergency head-CT turnaround (OP-23) at adopter-covered hospitals, post-2021

Treatment definition	coef.	(s.e.)	N	Treated hospitals
Any adopter-platform radiologist affiliated, 2019	−1.060	(0.951)	12,707	475
Adopter-majority coverage, 2019	−1.313	(1.244)	12,707	214
Radiology Partners present, 2019	−0.843	(1.079)	12,707	343
vRad present, 2019	−1.429	(1.451)	12,707	142

C. Hospitals whose coverage switched to an adopting platform: turnaround, protocoling, and ordering

	Post-switch	Before the switch (base −1)			Switching
Outcome (percentage points)	coef. (s.e.)	$\leqslant -3$	−2	N	hospitals
Head CT read within 45 minutes (OP-23), weighted	−3.76 (1.61)	+3.29 (2.15)	+1.04 (2.54)	11,831	193
Abdominal CT with and without contrast (OP-10)	0.05 (0.20)	+0.50 (0.30)	−0.01 (0.21)	27,868	302
Placebo: lumbar MRI for low back pain (OP-8)	−0.39 (0.66)	−2.01 (1.28)	−1.30 (1.20)	7,535	104
Placebo: cardiac imaging before surgery (OP-13)	−0.08 (0.13)	−0.64 (0.26)	−0.43 (0.26)	14,468	176

NOTE.—Panel A: log radiologists billing from the market, 2013–2024; market and year fixed effects and ln pop × clock; coverage from the 2017 hospital affiliations (145 vRad-covered markets, 118 covered by another teleradiology firm). Panel B: OP-23, the share of emergency head CTs read within 45 minutes, 2018–2025; hospital and period fixed effects. Panel C: hospitals that became adopter-majority in 2019, 2021 or 2023 (134, 104, 80) against never-covered hospitals; OP-23 is observed from 2018, so its earliest bin rests on the later switchers; OP-10 is a protocoling share, OP-8 and OP-13 shares of orders judged inappropriate (lower is better). Panels B and C: weighted by the measure's sample size (unweighted: −1.182 to −0.472 in panel B; −3.93 (1.45) in panel C).

Standard errors in parentheses, clustered by market (panel A) and by hospital (panels B and C); N: hospital-periods in panels B and C.

SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files; the Physician Compare and Doctors & Clinicians files; CMS Hospital Compare; the authors' record of documented AI adoption.

is +1.5 (2.0) points during the window and +4.9 (1.9) after it, the wrong pattern for an on–off response. The add-on paid for software use, so volume is a downstream outcome, and the window opened during the pandemic.

C.3 Billing Addresses versus Reading

Table C.3 reports the decomposition discussed in Subsection 5.3. The frozen-address headcount assigns each radiologist permanently to the market of first billing and counts those billing from anywhere; its gradient is the exit-and-entry component of the baseline, and the residual is readers who bill on from another market. Panel C characterizes those movers among the 2019 non-platform readers of small markets.

Table C.3: Billing addresses: current versus frozen, incumbent and entrant margins, and movers

	All markets		Fewer than 10 radiologists		10 or more	
Outcome	coef. (s.e.)	N	coef. (s.e.)	N	coef. (s.e.)	N
A. Log radiologists by billing address: exposure × clock						
Current address	−0.088 (0.013)	8,375	−0.124 (0.019)	4,266	0.061 (0.047)	3,918
Frozen first address	−0.029 (0.008)	8,414	−0.053 (0.011)	4,372	0.087 (0.034)	3,919
B. Incumbents, entrants and exit rates: exposure × clock						
Log 2013 incumbents billing	−0.097 (0.013)	7,603	−0.124 (0.020)	3,872		
ln(1+ post-2013 entrants)	0.080 (0.018)	6,967	0.014 (0.017)	3,535		
Entrants' share of headcount	0.011 (0.008)	6,967	−0.004 (0.012)	3,535		
Incumbent exit rate	−0.006 (0.005)	7,603	−0.001 (0.005)	3,872		
Entrant exit rate	0.007 (0.018)	3,714	−0.002 (0.013)	926		
Entry rate	0.006 (0.003)	7,714	0.005 (0.005)	3,955		

C. 2019 non-platform readers of markets with fewer than 10 radiologists, by where they billed in 2024

	Readers	Lists origin	Markets listed	Multi-market	Median Δ ln	To hub
Billing from the origin market	683	0.69	1.87	0.39	−0.35	0.00
Billing from another market	120	0.23	2.25	0.53	−0.38	0.68

NOTE.—Panel A: the frozen address assigns each radiologist to the market of first billing (2012 onward), so the count moves only through exit and entry; the current-address row reproduces Table 2. Panel B: Equation (9) with the size control, 2013–2024 (entrant rows 2014–2023); exit rates are the share not billing from the market the next year, the entry rate new NPIs over headcount. Panel C: readers still billing in 2024, by whether their 2024 market is the 2019 one; columns: share listing the origin market in 2025, mean markets listed, share listing more than one, median change in log services 2019–2024, share whose 2024 market has 50 or more radiologists. Standard errors in parentheses, clustered by market; N: market-years (panels A and B).
SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files; the Physician Compare and Doctors & Clinicians files.

C.4 On-Site Procedures

Table C.4 reports the procedure gradients by procedure group, setting and margin, for all markets and for markets with fewer than ten radiologists; Subsection 6.3 discusses them. The codes pay the billing physician for performing the procedure on the patient—placing a needle or catheter or conducting a fluoroscopic examination in real time—rather than for interpreting images afterwards, and are summed over their revisions. Radiologists bill most of them nationally (shares by group in

the table note); the versions of the aspiration, lumbar puncture and PICC codes without imaging, small since the code revisions of 2019 and 2020 (American Medical Association, 2018, 2019), are billed mostly by others. Ultrasound and fluoroscopic guidance codes are excluded: radiologists bill about one in seven of them, the rest being injections by anesthesiologists, orthopedic surgeons and pain specialists.

Table C.4: On-site procedures billed by radiologists and by others, by market size

	Zero cells	All markets		Markets with fewer than 10 radiologists				
	($n < 10$)	Log	N	Log	N	$\ln(1+n)$	LPM: any	N
A. Procedures billed by radiologists from the market								
All on-site procedures	29%	−0.186 (0.043)	7,407	−0.260 (0.076)	3,410	−0.418 (0.074)	−0.048 (0.016)	4,848
image-guided procedures	49%	−0.174 (0.054)	6,398	−0.230 (0.097)	2,457	−0.355 (0.070)	−0.061 (0.014)	4,848
fluoroscopic studies	35%	−0.219 (0.055)	7,032	−0.249 (0.096)	3,119	−0.418 (0.067)	−0.064 (0.016)	4,848
in the facility setting	31%	−0.209 (0.046)	7,299	−0.355 (0.078)	3,318	−0.426 (0.095)	−0.047 (0.021)	4,848
Radiologists billing any procedure	29%	−0.128 (0.019)	7,407	−0.172 (0.037)	3,410	−0.135 (0.018)	−0.048 (0.016)	4,848
B. The same procedures billed by others								
Other physicians and practitioners	65%	0.032 (0.037)	5,379	0.007 (0.061)	1,668	−0.105 (0.052)	−0.034 (0.013)	4,848
Anyone	23%	−0.132 (0.031)	7,823	−0.192 (0.056)	3,727	−0.353 (0.070)	−0.040 (0.013)	4,848

C. Where procedures are physically done: payment localities, 2016–2024

	Done in locality	Billed from it	Difference	N
All on-site procedures	−0.053 (0.021)	−0.101 (0.029)	0.048 (0.015)	756
image-guided procedures	−0.008 (0.025)	−0.060 (0.033)	0.052 (0.017)	747
fluoroscopic studies	−0.129 (0.028)	−0.167 (0.041)	0.039 (0.021)	756

D. Reassigning procedures to listed locations: markets with fewer than 10 radiologists, $\ln(1+n)$

	coef. (s.e.)	N
At the claims billing address	−0.458 (0.096)	2,070
Reassigned to listed practice locations	−0.189 (0.074)	2,070
and to hospital affiliations (at most three)	0.040 (0.057)	2,070

NOTE.—Coefficients on exposure × clock from Equation (9) with the size control, 2013–2024; the log columns use market-years with a positive count, the others the balanced panel; zero cells: share of market-years with no such procedure in markets with fewer than 10 radiologists; procedures (Appendix Table A.1) at the billing provider's practice address. Panel C: as Table 3, panel D, for the procedure basket; done in locality from the Physician/Supplier Procedure Summary. Panel D: claims years 2017, 2019, 2021, 2023 and 2024; procedures billed from outside a small market are assigned to it, split equally, where the radiologist lists a practice location there (last row: or a hospital affiliation). Radiologists' 2024 share of these services nationally: 86–96 percent for biopsies, CT guidance, drainage and fluoroscopic studies; 44–64 percent for aspiration, lumbar puncture, paracentesis, thoracentesis and PICC placement; 13–14 percent for the excluded guidance codes.
Standard errors in parentheses, clustered by market (panel C: by locality); N: market-years (panel C: locality-years).
SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files; the Physician/Supplier Procedure Summary; the CMS ZIP-to-locality files; the Physician Compare and Doctors & Clinicians files.

C.5 The Interstate Medical Licensure Compact

States became operational in the compact at dates spread from 2017 to 2026; 138 markets are in states that have not joined. Panels A and B of Table C.5 report the staggered-entry estimates.

Operational status moves neither billing headcount nor the organizational shares; the one margin it moves is the count of readers listing a single market, which falls, weakly consistent with the compact easing multi-state listing for readers already in place; the direct measure, the share listing more than one market, does not move. Family-practice and psychiatry headcounts show no compact effect.

Table C.5: The Interstate Medical Licensure Compact

Outcome and sample	coef.	(s.e.)	N
A. Billing headcount, state operational in the compact			
All markets	0.016	(0.017)	8,375
Markets with fewer than 10 radiologists	0.033	(0.022)	4,266
10–49	−0.027	(0.024)	2,610
50 or more	−0.015	(0.023)	1,308
Compact × exposure (all markets)	−0.036	(0.029)	8,375
B. Directory margins and placebo professions			
Share of listed radiologists in groups spanning 5 or more states	−0.007	(0.024)	4,122
Share listing locations in more than one market	0.008	(0.019)	4,161
Log radiologists listing a single market	−0.068	(0.033)	4,161
Family-practice physicians, billing headcount	−0.003	(0.012)	10,698
Psychiatrists, billing headcount	0.019	(0.021)	8,505

NOTE.—Event study of billing headcount around the state's operational year (base −1): coefficients between −0.005 and +0.057, none with $|t| > 1.5$. Treatment: the market's state is operational in the compact; market and year fixed effects with exposure × clock and ln pop × clock, estimation sample of Table 2. Panel B: the five directory snapshots with a post-2021 indicator in place of the clock; the two placebo rows are annual headcount regressions, 2013–2024.
Standard errors in parentheses, clustered by state.
SOURCES.—Authors' calculations from the Medicare Physician & Other Practitioners files; the Physician Compare and Doctors & Clinicians files; the Interstate Medical Licensure Compact Commission; Census population estimates.

C.6 Adopter Entry versus Non-Adopter Entry

A group's entry into a market is dated to the first Physician Compare snapshot in which one of its affiliated radiologists lists a location there; markets where the group was already present in 2017 are excluded from the corresponding design. Documented adopters are the groups linked to the adoption record, each from the year of its first interpretive-AI deployment (triage, detection or second read; report-drafting and operational tools do not count), and markets where a group that later adopted was present in 2017 are excluded; non-adopting multi-state groups are those spanning five or more states with no documented adoption of any kind. Table C.6 reports the post-entry indicators alone and interacted with exposure, by market class. Entry moves headcount for neither group on average; in small markets the exposure gradient is larger after adopter entry than after non-adopter entry, though not significantly so, and in markets with ten to forty-nine radiologists the adopter interaction is positive (0.183, s.e. 0.098).

Table C.6: Adopter entry versus non-adopter entry: billing headcount

	Documented adopter		Non-adopting multi-state group	
	coef. (s.e.)	N	coef. (s.e.)	N
All markets				
Post-entry indicator, entered alone	0.034 (0.025)	7,009	−0.019 (0.025)	5,960
Post-entry indicator, with the interaction below	0.048 (0.043)	5,624	−0.030 (0.032)	5,624
Post-entry × exposure	−0.061 (0.068)	5,624	0.006 (0.042)	5,624
Fewer than 10 radiologists				
Post-entry indicator, entered alone	0.052 (0.046)	3,992	−0.010 (0.038)	3,649
Post-entry indicator, with the interaction below	0.181 (0.084)	3,505	−0.030 (0.058)	3,505
Post-entry × exposure	−0.239 (0.113)	3,505	0.015 (0.067)	3,505
10–49				
Post-entry indicator, entered alone	0.024 (0.028)	2,262	−0.053 (0.027)	1,854
Post-entry indicator, with the interaction below	−0.091 (0.049)	1,722	−0.011 (0.043)	1,722
Post-entry × exposure	0.183 (0.098)	1,722	−0.088 (0.075)	1,722

NOTE.—Adopter entry is the first snapshot in which a group that has deployed interpretive AI (triage, detection or second read, dated from the adoption record) lists a location in the market, non-adopter entry the first in which a group spanning five or more states and absent from the record does; markets where such a group, or a group that later adopted, was present in 2017 are excluded; same fixed effects and controls; both entries enter the same regression in the interaction rows. The 50-or-more class is omitted: too few groups enter after 2017 to estimate the interaction. The adopter-entry event study is flat: coefficients from −0.04 to +0.04.
Standard errors in parentheses, clustered by market; N: market-years.
SOURCES.—Authors' calculations from the Physician Compare and Doctors & Clinicians files; the authors' record of documented AI adoption; the Medicare Physician & Other Practitioners files.

D Quantitative Model and Counterfactuals

This appendix gives the construction behind Section 7: the market-level model, the data and calibration, the untargeted moments, the counterfactual formulas and the sensitivity analysis. Nothing here re-estimates the causal effects of the main text.

D.1 Market-Level Quantitative Model

Subsection 7.1 defines the market-level objects: excess capacity π_m from the equal allocation of V_m^0 across the J_m hospitals of market m, with $n(v) = \max\{n_{\min}, \lceil v/\bar{h} \rceil\}$; the cross-market propensity $q_m^t = \Lambda(\alpha_t + \lambda\pi_m)$; and the retained share $r_m = \exp\{-\psi(q_m^1 - q_m^0)\}$. Potential production is $D_m^1 = (1+g_m)V_m^0$. Before accounting for practice viability, locally retained production is $r_m D_m^1$, while $(1-r_m)D_m^1$ is reorganized across markets. A local production unit remains viable if its retained volume per hospital satisfies $r_m D_m^1/J_m \geqslant \theta\bar{h}$, where $\theta = c_L/p$ and $n_{\min} = 1$. When this condition fails, its remaining interpretations are also transferred to pooled production. We denote the resulting final allocation by V_{Im}^1 and V_{Rm}^1, with $V_{Im}^1 + V_{Rm}^1 = D_m^1$. Reorganized interpretations are assigned to markets with at least 25 radiologists in proportion to their 2013 headcounts. This appendix adds the capacity accounting.

To translate production into reader-capacity requirements, let $a \geqslant 1$ denote task productivity, normalized to $a = 1$ in the benchmark. Initial requirements are

$$L^0 = \sum_m \frac{V_m^0}{\bar{h}}(1+\pi_m).$$

For post-transition production, we use the baseline excess-capacity penalty for interpretations retained in their original local units and assume that reorganized interpretations can be pooled without this penalty:

$$L^1 = \sum_m \left[\frac{V_{Im}^1}{a\bar{h}}(1+\pi_m) + \frac{V_{Rm}^1}{a\bar{h}} \right].$$

This is a fixed-penalty accounting approximation: it applies initial excess-capacity ratios to retained production rather than recomputing integer staffing requirements after each volume adjustment. It isolates the capacity savings associated with pooling under the calibrated geographic allocation.

The capacity required to serve post-transition demand entirely within the original local units is $L_{\text{local}}^1 = \sum_m D_m^1(1+\pi_m)/\bar{h}$. With $a = 1$, reorganization therefore releases

$$L_{\text{local}}^1 - L^1 = \sum_m \frac{V_{Rm}^1 \pi_m}{\bar{h}},$$

giving a proportional capacity saving

$$S_V = 1 - \frac{L^1}{L_{\text{local}}^1}.$$

In the benchmark, $S_V = 2.5$ percent. This measures the reader capacity saved through pooling relative to the specified local-production benchmark, not the causal effect of AI verification. It also provides the normalization for the equal-saving task-productivity counterfactual.

D.2 Data Construction and Calibration

The four organizational parameters $(\lambda, \alpha_0, \alpha_1, \psi)$ are calibrated sequentially to observed affiliation and billing patterns.

Baseline volume. V_m^0 consists of imaging interpretations (procedure codes 70010–79999, excluding the computer-aided-detection add-ons, and the pre-2018 mammography G-codes) billed in 2013 by radiologists whose billing ZIP lies in market m, including both facility and office settings. This yields 81 million interpretations, of which 69 million are in facility settings.

Reader capacity. We set $\bar{h} = 3{,}984$, the 75th percentile of annual interpretations per radiologist across all settings in 2013. This is a reader-equivalent capacity normalization, not an estimate of the technological maximum.

Production units. The public files do not identify the hospital at which each image was produced, so hospital-level interpretation volumes cannot be recovered. The benchmark allocates V_m^0 equally across the J_m hospitals recorded in the 2017 directory for market m; Subsection D.5 repeats the calibration on facility-setting interpretations alone, with alternative production units and with higher minimum staffing.

Cross-market organization. The 2017 and 2025 hospital-directory snapshots identify affiliated radiologists and their practice CBSAs. An affiliation is cross-market when the hospital and radiologist are located in different CBSAs. A fractional logit of 2017 market shares on π_m, weighted by affiliation counts, disciplines λ and α_0; α_1 matches the aggregate 2025 affiliation share.

Demand growth. We construct market demand from imaging events per 1,000 Medicare fee-for-service beneficiaries, classified by beneficiary residence, and the corresponding number of fee-for-service beneficiaries (Medicare Monthly Enrollment, July):

$$D_{mt} = \frac{\text{Imaging events per 1,000}_{mt}}{1{,}000} \times \text{FFS beneficiaries}_{mt}.$$

Medicare Advantage penetration increased from 29 to 51 percent over the period. Market-specific growth is $g_m = D_{m,2024}/D_{m,2013} - 1$, giving potential post-transition production $D_m^1 = (1 + g_m)V_m^0$.

Reallocation intensity and viability. The viability threshold is fixed at $\theta = 1$, rather than fitted to the geographic gradient. Allowing the threshold to fit the retention ratio for markets with fewer than ten radiologists yields $\theta = 1.26$ (Table D.2, last column).

D.3 Targeted and Untargeted Moments

Table D.1 summarizes the benchmark inputs and calibration. The parameters λ and α_0 are disciplined by the 2017 cross-sectional relationship between excess capacity and cross-market affiliations, α_1 matches the aggregate 2025 affiliation share, and ψ matches the 2024-to-2013 billing ratio for markets with fewer than 25 radiologists.

Table D.1: Benchmark quantitative calibration

Object	Value	Discipline
A. Direct inputs		
Reader capacity, $\bar{h}$	3,984	2013 75th percentile of interpretations per billing radiologist, all settings
Minimum staffing, $n_{\min}$	1	benchmark normalization
Practice viability, c_L/p	1	the fee covers one fully utilized reader; a practice whose retained volume per hospital falls below one reader-load exits
Task productivity, a	1	verification-only benchmark
B. Organizational calibration		
Sensitivity to excess capacity, λ	0.51	2017 cross-section of market affiliation shares
Baseline intercept, α_0	-1.23	2017 cross-section
Post-AI intercept, α_1	-0.75	2025 aggregate cross-market affiliation share
Reallocation intensity, ψ	3.10	2024/2013 ratio of reads billed from markets under 25 radiologists
C. Production units by market size (median reader-loads per hospital; share of hospitals below one)		
Fewer than 10 radiologists	1.8; 30%	2013 interpretations split equally across the market's 2017 hospitals, divided by $\bar{h}$
10–24	5.0; 1%	
25–49	6.1; 0%	
50 or more	7.9; 0%	

NOTE.—Market interpretation volumes, hospital counts and fee-for-service demand growth are observed inputs. The market-size ratios above 25 radiologists and the aggregate ratio are not targeted.
SOURCE.—Authors' calculations from the calibrated model of Section 7.

The untargeted moments—the ratios by market-size class, aggregate billing, and the share of initial local reader positions in units that fall below the viability threshold—are in Table 9 and discussed in Subsection 7.1; the model-implied positions are positions in nonviable units, not observed individual exits.

D.4 Counterfactual Implications

The counterfactual exercises vary the production or demand assumptions relative to the benchmark. Organizational parameters are held fixed unless recalibration is explicitly indicated.

Removing production indivisibilities. Setting $\pi_m = 0$ for every market and $\theta = 0$, and recomputing q_m^0, q_m^1, and r_m, raises the predicted 2024-to-2013 billing ratio for markets with fewer than ten radiologists from 0.55 to 0.65 (or 0.55 when only $\pi_m = 0$; Table 10). Relative to divisible

production, which requires $L^0_{\text{div}} = \sum_m V^0_m/\bar{h}$ readers, discrete local staffing requires

$$E^0 = \frac{L^0}{L^0_{\text{div}}} - 1 \simeq 8.0\%$$

more capacity in 2013. The benchmark reorganization subsequently releases 2.5 percent of the capacity required to serve post-transition demand within the original local units, as measured by S_V in Subsection D.1.

Organizational reallocation versus task productivity. Consider a task-productivity improvement a_P that reduces reader requirements while holding the geographic allocation of production fixed. With no change in organizational boundaries, its proportional capacity saving is $1 - 1/a_P$. Equating this to the benchmark saving from reorganization gives

$$1 - \frac{1}{a_P} = S_V, \qquad a_P = \frac{1}{1 - S_V} \simeq 1.025.$$

By construction, the two counterfactuals generate the same aggregate capacity saving, and both retain aggregate production at approximately 0.99 of its 2013 level. Their geographic implications differ because the task-productivity counterfactual holds production locations fixed (Table 10). This restriction isolates the organizational margin; it does not imply that productivity improvements could not themselves induce outsourcing or practice exit.

Demand growth. Let g denote a common proportional increase in demand, and let $L^j(g)$ be the aggregate reader-capacity requirement under counterfactual j. The break-even demand-growth rate satisfies $L^j(g^*_j) = L^0$. Holding the calibrated allocation and capacity-saving factors fixed, the organizational reallocation has a break-even rate of approximately $g^*_V = 2.6$ percent. Re-solving the threshold with practice viability recomputed at each growth rate changes it by less than 0.1 percentage point (Table 10, panel C). The equal-saving productivity benchmark has the same threshold by construction. Applying both savings multiplicatively gives

$$g^*_{V+P} = \frac{1}{(1 - S_V)^2} - 1 \simeq 5.2\%.$$

These thresholds are accounting comparisons conditional on the specified allocation and staffing assumptions, rather than predictions from a model that reoptimizes organizational boundaries as demand changes.

D.5 Alternative Calibrations and Sensitivity

Table D.2 examines alternative assumptions about production indivisibilities and practice viability. We recalibrate the organizational parameters to the same moments under minimum staffing requirements $n_{\min} \in \{2, 3, 4\}$, when the CBSA is treated as a single production unit, and under

viability thresholds $\theta \in \{0, 0.75, 1.25\}$. In the single-unit specification, excess capacity is

$$\pi_m^{CBSA} = \frac{\max\{n_{\min}, \lceil V_m^0/\bar{h} \rceil\}\bar{h}}{V_m^0} - 1.$$

Increasing the minimum staffing complement raises measured excess capacity in the smallest markets from 0.21 under $n_{\min} = 1$ to 0.33, 0.55, and 0.83 under $n_{\min} = 2, 3, 4$, respectively. The 2017 affiliation cross-section then implies a smaller λ (Table D.2, panel C), while the predicted billing ratios remain 0.54, 0.68, 0.92, and 1.10 across the four market-size classes. Treating the CBSA as a single production unit reduces measured excess capacity, with the same predicted ratios after recalibration. Thus, the geographic gradient is stable across these alternative definitions of production indivisibilities, conditional on matching the same organizational moments. The benchmark retains the conservative assumption $n_{\min} = 1$.

The viability threshold matters primarily in the smallest markets. For markets with fewer than ten radiologists, the predicted billing ratio is 0.60 when $\theta = 0$, 0.57 when $\theta = 0.75$, and 0.53 when $\theta = 1.25$, compared with 0.55 in the benchmark and 0.53 in the data. The predicted ratios in the remaining market-size classes change by at most 0.04. Alternative rules for allocating reorganized reads across hub markets move only the two larger classes, between 0.68 and 1.01 for markets with 25–49 radiologists and between 1.09 and 1.14 for markets with fifty or more, and leave the targeted classes and the capacity saving unchanged (panel D of the same table).

Table D.3 compares the benchmark with a specification using facility-setting interpretations alone, recalibrating reader capacity and all organizational parameters. The facility-only specification closely reproduces the small-market gradient, while yielding somewhat greater initial excess capacity (9.0 versus 8.0 percent) and production reallocation (32.8 versus 26.7 percent). Aggregate capacity savings range from 0.7 to 3.3 percent across the two production measures and alternative staffing calculations. The facility-only specification fits aggregate billing less closely, predicting a ratio of 0.98 against 0.85 in the data, consistent with changes in the distribution of imaging between facility and office settings.

Table D.2: Sensitivity of the quantitative calibration

	Benchmark	Minimum staffing			CBSA as	Practice viability c_L/p			
	$n_{\min} = 1$	$n_{\min} = 2$	$n_{\min} = 3$	$n_{\min} = 4$	one unit	0	0.75	1.25	fitted
A. Excess capacity in 2013, volume-weighted, by 2013 market size									
Fewer than 10 radiologists	0.21	0.33	0.55	0.83	0.13	0.21	0.21	0.21	0.21
10–24	0.08	0.08	0.10	0.13	0.03	0.08	0.08	0.08	0.08
25–49	0.07	0.07	0.08	0.09	0.02	0.07	0.07	0.07	0.07
50 or more	0.07	0.07	0.07	0.07	0.00	0.07	0.07	0.07	0.07
B. Model-implied reads billed from the market, 2024/2013									
Fewer than 10 radiologists	0.55	0.54	0.54	0.54	0.54	0.60	0.57	0.53	0.53
10–24	0.68	0.68	0.68	0.69	0.68	0.65	0.67	0.69	0.69
25–49	0.92	0.92	0.92	0.92	0.92	0.91	0.92	0.92	0.92
50 or more	1.10	1.10	1.10	1.10	1.10	1.10	1.10	1.10	1.10
C. Recalibrated parameters									
λ	0.51	0.24	0.20	0.17	0.70	0.51	0.51	0.51	0.51
ψ	3.10	3.09	3.07	3.05	3.07	3.78	3.41	2.86	2.83
c_L/p	1.00	1.00	1.00	1.00	1.00	0.00	0.75	1.25	1.26
S_V (percent)	2.5	3.4	4.8	6.3	0.8	2.3	2.5	2.4	2.4

D. Allocation of reorganized reads across hub markets, parameters fixed

Allocation rule	Reads billed from the market, 2024/2013				S_V (%)
	Fewer than 10	10–24	25–49	50 or more	
markets with ⩾25, by 2013 headcount (benchmark)	0.55	0.68	0.92	1.10	2.5
markets with ⩾25, by 2013 volume	0.55	0.68	1.01	1.09	2.5
markets with ⩾50 only, by 2013 headcount	0.55	0.68	0.68	1.14	2.5
markets with ⩾25, by headcount × demand growth	0.55	0.68	0.90	1.11	2.5

NOTE.—Each column recalibrates λ, α_0, α_1 and ψ to the benchmark's targets after changing the construction of excess capacity (the minimum complement per hospital, or the whole CBSA as one production unit) or the practice-viability threshold (the benchmark is one reader-load; 0 removes the exit margin). The market-size ratios are not targeted. The last column fits c_L/p to the ratio for markets with fewer than 10 radiologists. Panel C reports the recalibrated parameters; S_V is the aggregate reader-capacity saving from reorganization relative to serving 2024 demand within the original local units, in percent. Panel D holds the benchmark parameters fixed and changes only where reorganized reads are billed.
SOURCE.—Authors' calculations from the calibrated model of Section 7.

Table D.3: Alternative production measures and capacity accounting

	All interpretations (benchmark)	Facility only
Annual reader capacity, $\bar{h}$	3,984	3,779
Initial excess capacity (%)	8.0	9.0
Production reorganized (%)	26.7	32.8
Capacity saving, fixed baseline slack (%)	2.5	3.3
Capacity saving, recomputed staffing (%)	1.5	0.7
Positions in exiting units, fewer than 10 radiologists (%)	21	23
Billing ratio, 2024/2013: observed / model		
Fewer than 10 radiologists	0.53 / 0.55	0.49 / 0.49
10–24	0.69 / 0.68	0.62 / 0.61
25–49	0.87 / 0.92	0.78 / 0.90
50 or more	1.09 / 1.10	0.94 / 1.11
Aggregate	0.98 / 0.99	0.85 / 0.98

NOTE.—The benchmark includes interpretations in both facility and office settings; the alternative includes facility-setting interpretations only. Reader capacity and the organizational parameters are recalibrated for the alternative specification. Capacity savings are measured relative to the reader requirements for serving post-transition demand within the original local units. The fixed-slack calculation retains the 2013 excess-capacity penalty for locally billed production; the alternative recomputes whole-reader staffing after reallocation. Positions in exiting units are a percentage of initial local reader positions in markets with fewer than ten radiologists. Billing ratios compare 2024 with 2013 and report observed values followed by model predictions.
SOURCE.—Authors' calculations from the calibrated model of Section 7.

E Data and Documentary Sources

American College of Radiology (2025). ACR practice parameter for communication of diagnostic imaging findings. Revised 2025 (Resolution 9). Reston, VA: American College of Radiology.

American Medical Association (2018). *CPT 2019 Professional Edition*. Chicago: American Medical Association.

American Medical Association (2019). *CPT Professional 2020*. Chicago: American Medical Association.

Centers for Medicare & Medicaid Services (2012). Revised and clarified place of service (pos) coding instructions. Medicare Claims Processing Manual Transmittal 2407, Change Request 7631, 3 February 2012.

Centers for Medicare & Medicaid Services (2020a). Clinical laboratory improvement amendments (CLIA) laboratory guidance during COVID-19 public health emergency. Memorandum QSO-20-21-CLIA, 26 March 2020.

Centers for Medicare & Medicaid Services (2020b). Medicare program; hospital inpatient prospective payment systems for acute care hospitals and the long-term care hospital prospective payment system and final policy changes and fiscal year 2021 rates; final rule. 85 Federal Register 58432, 18 September 2020 (new technology add-on payment for AI detection of large-vessel occlusion).

Centers for Medicare & Medicaid Services (2021). Medicare program; hospital inpatient prospective payment systems for acute care hospitals and the long-term care hospital prospective payment system and policy changes and fiscal year 2022 rates; final rule. 86 Federal Register 44774, 13 August 2021 (one-year extension of the add-on payment for ContaCT).

Centers for Medicare & Medicaid Services (2023). Clinical laboratory improvement amendments of 1988 (clia) post-public health emergency (phe) guidance. Memorandum QSO-23-15-CLIA, 11 May 2023, revised 23 September 2025.

Centers for Medicare & Medicaid Services (2024). ZIP code to carrier locality file, 2016–2024 releases. Data set, Medicare Physician Fee Schedule. https://www.cms.gov/medicare/payment/fee-schedules.

Centers for Medicare & Medicaid Services (2025a). Medicare claims processing manual, chapter 13: Radiology services and other diagnostic procedures, section 20.1. CMS Pub. 100-04, Rev. 13150, 11 April 2025. https://www.cms.gov/regulations-and-guidance/guidance/manuals/downloads/clm104c13.pdf.

Centers for Medicare & Medicaid Services (2025b). Physician compare national downloadable file (2017, 2019) and doctors and clinicians national downloadable file (2021, 2023, 2025). https://data.cms.gov/provider-data/topics/doctors-clinicians; archived releases at https://data.cms.gov/provider-data/archived-data/doctors-clinicians.

Centers for Medicare & Medicaid Services (2026a). Hospital compare: Outpatient imaging efficiency measures (OP-8, OP-10, OP-13) and timely and effective care measure OP-23. Provider Data Catalog. https://data.cms.gov/provider-data/topics/hospitals.

Centers for Medicare & Medicaid Services (2026b). Hospital general information. Provider Data Catalog. https://data.cms.gov/provider-data/dataset/xubh-q36u.

Centers for Medicare & Medicaid Services (2026c). Medicare claims processing manual, chapter 1: General billing requirements, section 10.1.1. CMS Pub. 100-04, Rev. 13725, 10 April 2026.

Centers for Medicare & Medicaid Services (2026d). Medicare geographic variation public use file, by national, state and county: Historical file 2007–2013 ("2007–2013 medicare ffs geographic variation public use file", may 2024 release) and current series 2014–2024. https://data.cms.gov/summary-statistics-on-use-and-payments/medicare-geographic-comparisons/medicare-geographic-variation-by-national-state-county.

Centers for Medicare & Medicaid Services (2026e). Medicare inpatient hospitals—by provider and service. https://data.cms.gov/provider-summary-by-type-of-service/medicare-inpatient-hospitals.

Centers for Medicare & Medicaid Services (2026f). Medicare monthly enrollment. https://data.cms.gov/summary-statistics-on-beneficiary-enrollment/medicare-and-medicaid-reports/medicare-monthly-enrollment.

Centers for Medicare & Medicaid Services (2026g). Medicare physician & other practitioners—by provider and service, and by provider, calendar years 2013–2024. Public use files. https://data.cms.gov/provider-summary-by-type-of-service/medicare-physician-other-practitioners.

Centers for Medicare & Medicaid Services (2026h). Physician fee schedule. 42 CFR Part 414, Subpart B; annual final rules.

Centers for Medicare & Medicaid Services (2026i). Physician/supplier procedure summary, calendar years 2013–2024. Data set. https://data.cms.gov/summary-statistics-on-use-and-payments/physiciansupplier-procedure-summary.

Centers for Medicare & Medicaid Services (2026j). Rural emergency hospitals. https://www.cms.gov/medicare/health-safety-standards/certification-compliance/rural-emergency-hospitals, accessed October 2026.

Code of Federal Regulations (2026a). Title 21, part 892, radiology devices, sections 892.2050, 892.2060, 892.2070 and 892.2080. Electronic Code of Federal Regulations.

Code of Federal Regulations (2026b). Title 42, part 493, laboratory requirements, sections 493.35, 493.43 and 493.55. Electronic Code of Federal Regulations.

Code of Federal Regulations (2026c). Title 42, section 415.120, conditions for payment: Radiology services. Electronic Code of Federal Regulations.

Interstate Medical Licensure Compact Commission (2024). Participating states: Compact state map and implementation dates. https://www.imlcc.com/participating-states/, accessed August 2026.

MEDNAX, Inc. (2015). MEDNAX completes acquisition of vRad. Press release, Form 8-K Exhibit 99.1, 18 May 2015, SEC accession 0001193125-15-190959.

MEDNAX, Inc. (2017). Annual report on form 10-k for the fiscal year ended december 31, 2016. Filed 10 February 2017, SEC accession 0001193125-17-038250.

MEDNAX, Inc. (2020). Form 8-k, exhibit 99.2: MEDNAX completes sale of MEDNAX radiology solutions to radiology partners. 21 December 2020, SEC accession 0001193125-20-323140.

National Bureau of Economic Research (2026). Medicare provider utilization and payment data: Physician and other supplier, calendar year 2012, aggregate (by-npi) file. NBER CMS data mirror, data.nber.org, http://data.nber.org/providerchargepuf/2012/. File obtained 18 August 2026 and deposited with its checksum in the replication package.

Radiology Partners (2021). Radiology partners and aidoc to accelerate the use of AI as the standard of care in radiology. Press release, 6 April 2021, https://www.radpartners.com/2021/04/radiology-partners-and-aidoc-to-accelerate-the-use-of-ai-as-the-standard-of-care-in-radiology/.

Radiology Partners (2023). Radiology partners deploys clinical AI across more than 20 million annual patient exams. Press release, 13 November 2023, https://www.radpartners.com/2023/11/radiology-partners-deploys-clinical-ai-across-more-than-20-million-annual-patient-exams/.

Strategic Radiology (2026). About strategic radiology. https://www.strategicradiology.org/, accessed October 2026.

U.S. Bureau of Labor Statistics (2026). Occupational employment and wage statistics, metropolitan and nonmetropolitan area estimates, 2011–2025. https://www.bls.gov/oes/tables.htm.

U.S. Census Bureau (2017). 2017 gazetteer files: Core based statistical areas. https://www.census.gov/geographies/reference-files/time-series/geo/gazetteer-files.html.

U.S. Census Bureau (2025). Annual estimates of the resident population for metropolitan and micropolitan statistical areas, vintages 2019 and 2024. Tables CBSA-MET-EST2019-ANNRES, CBSA-MIC-EST2019-ANNRES, CBSA-MET-EST2024-POP and CBSA-MIC-EST2024-POP. https://www.census.gov/data/tables/time-series/demo/popest/2020s-total-metro-and-micro-statistical-areas.html.

U.S. Department of Housing and Urban Development (2017). HUD USPS ZIP code crosswalk files: ZIP–CBSA, fourth quarter 2017. https://www.huduser.gov/portal/datasets/usps_crosswalk.html.

U.S. Food and Drug Administration (2017). De novo classification request den160056: Philips intellisite pathology solution. Decision of 12 April 2017.

U.S. Food and Drug Administration (2026). List of artificial intelligence-enabled medical devices. Device list, updated periodically. https://www.fda.gov/medical-devices/artificial-intelligence-enabled-medical-devices/list-artificial-intelligence-enabled-medical-devices.

U.S. Office of Management and Budget and U.S. Census Bureau (2018). Core based statistical areas, metropolitan divisions, and combined statistical areas: September 2018 delineation file. Data set. https://www2.census.gov/programs-surveys/metro-micro/geographies/reference-files/2018/delineation-files/list1_Sep_2018.xls.

vRad (Virtual Radiologic) (2020). New radiology AI models reduce time to care. vRad blog, 1 July 2020, https://blog.vrad.com/new-ai-models-reduce-time-to-care.

vRad (Virtual Radiologic) (2021). vRad quality assurance process. Document 67.E.BS.0821, August 2021, https://info.vrad.com/hubfs/PDFs/vRad_Quality_Assurance_Process.pdf.