Projected Population Health Impact of Autonomous AI Primary Care in US Physician-Shortage Counties

Alon Bergman, PhD[1,2] and Dan Zeltzer, PhD[3]

[1]Department of Medical Ethics & Health Policy, University of Pennsylvania, Philadelphia

[2]Department of Health Care Management, The Wharton School, University of Pennsylvania, Philadelphia

[3] Division of Health Policy and Management, UC Berkeley School of Public Health

Presenter:

Alon Bergman, PhD, 3600 Civic Center Blvd, Philadelphia, PA 19104

(Alon.Bergman@pennmedicine.upenn.edu).

Manuscript word count: 2,643

Abstract

More than 100 million Americans live in federally designated primary care shortage areas. Recent advances in large language models have made autonomous AI primary care a near-term possibility, yet the population-health consequences of national deployment remain unquantified. This paper develops a county-level decision-analytic model that combines 2025 Area Health Resources File physician supply data, CDC WONDER mortality data, and causal parameters from a prespecified systematic review of quasi-experimental studies of the primary care supply–outcome relationship. Across the 2,303 counties (74.1 percent of the sample; 141.9 million residents) falling below a 1:1,500 adequacy threshold, deployment under a moderate adoption scenario calibrated to observed telehealth utilization is projected to avert approximately 10,700 deaths and 5,400 preventable hospitalizations annually, a mortality reduction roughly twice that associated with ACA Medicaid expansion among low-income older adults. Benefits are concentrated in the South and in nonmetropolitan counties. Estimates remain policy-relevant under physician-equivalence factors as low as 25 percent. A bottom-up cost triangulation implies annual deployment costs of roughly $4–7 billion, or about $470,000 per death averted. The magnitudes are large enough, and robust enough to conservative assumptions about AI effectiveness, that autonomous primary care warrants evaluation as a national supply-side policy instrument rather than a clinical curiosity. Those returns, however, are not automatic: they accrue only where cost-sharing and implementation design reach the uninsured, rural, and low-connectivity populations in which the shortage, and therefore the projected benefit, is concentrated.

Introduction

More than 100 million Americans live in 1 of approximately 7,000 federally designated Health Professional Shortage Areas (HPSAs), and the Health Resources and Services Administration

(HRSA) projects a national primary care physician shortfall exceeding 70,000 by 2038.[1,2] Geographic maldistribution is the dominant driver of this shortage. Rural communities face physician-to-population ratios substantially below national averages, and more than 1,000 rural family physicians left the workforce between 2019 and 2024.[3] The consequences include higher rates of preventable mortality and ambulatory care–sensitive hospitalizations that track the geography of shortage closely.[4]

Recent advances in large language models (LLMs) have made autonomous AI primary care realistic in the near term. Frontier AI systems have been shown in controlled evaluations to match or exceed practicing physicians on medical knowledge benchmarks, structured diagnostic vignettes, objective structured clinical examinations (OSCEs), and, in 1 randomized double-blind evaluation, most quality and safety measures.[5,6,7] These capabilities align closely with the core cognitive tasks of primary care—history-taking, diagnosis, guideline-concordant management, and chronic disease monitoring.

Whether deploying an autonomous AI primary care physician (AI PCP) at national scale would meaningfully reduce the population health burden of these shortages has not been quantified. Using county-level physician supply data, published causal estimates of the supply–outcome relationship, and adoption rates grounded in observed telehealth utilization, we simulated the annual deaths and preventable hospitalizations that a nationally available AI PCP would avert under realistic adoption scenarios. These findings are intended to inform the regulatory and payment frameworks now being developed for autonomous clinical AI.

Methods

We conducted a county-level simulation study (decision analytical model) to estimate the annual population health impact of deploying an autonomous AI PCP in US counties with primary care physician supply below a predefined adequacy threshold. Outcomes were annual all-cause

deaths averted and Prevention Quality Indicator (PQI-90) preventable hospitalizations averted. We modeled a single steady-state year using 2025 county-level primary care supply, contemporaneous mortality and population data, and fixed causal parameters from the literature. Projections represent the annual benefit in the year after hypothetical full deployment. Because the simulation used only publicly available aggregate county-level data, it was classified as non–human subjects research and did not require institutional review board approval or informed consent.

Data Sources

We obtained county-level primary care physician supply from the 2025 Area Health Resources File.[8] All-cause mortality rates and population denominators came from the Centers for Disease Control and Prevention WONDER Multiple Cause of Death data (2018-2024).[9] HPSA designations came from the HRSA HPSA Shortage Designation file,[1] and rural-urban status was assigned using US Department of Agriculture (USDA) Rural-Urban Continuum Codes (RUCCs; nonmetropolitan, RUCC ≥4).[10] The analytic sample included 3,106 US counties in the 50 states and the District of Columbia after excluding territories and 47 counties lacking usable mortality data.

Simulation Framework

The simulation required a supply threshold q* to identify counties with a primary care supply gap that an AI PCP could address. Counties below q* constituted the treated population; those at or above q* were classified as adequately supplied and received no simulated benefit. We set q* at 66.7 physicians per 100,000 population (1:1,500 ratio), the alternative shortage-area threshold recommended by the HRSA Negotiated Rulemaking Committee in 2011 and used as the principal scenario in Basu et al.[11] This threshold falls between the median and 75th percentile of county primary care physician supply in the sample.

The AI PCP was modeled as deployed nationally via a digital platform to patients in counties below q*, with no additional eligibility restrictions. In treated counties, the effective physician-equivalent supply added was the product of the patient adoption rate τ and the local supply gap (q* minus observed supply). Counties at or above q* were unaffected. The model assumes AI PCP clinical effectiveness equivalent to that of an additional primary care physician in the source studies. We evaluated 3 prespecified adoption scenarios grounded in observed telehealth utilization: conservative (τ=0.19; rural adults), moderate (τ=0.31; nonrural adults), and upper (τ=0.47; peak Medicare primary care telehealth share in April 2020).[12,13]

Outcome Parameters

Outcome parameters were drawn from a prespecified systematic review of causally identified studies of the primary care supply–health outcome relationship (eAppendix A). We searched PubMed/MEDLINE, EconLit, Web of Science, Scopus, and NBER/SSRN for studies using instrumental variables, difference-in-differences with event-study validation, or patient-mover designs. Retained estimates (tier 1, strong identification; tier 2, moderate) were harmonized to deaths or admissions per 100,000 population per additional primary care provider per 10,000 population.

Mortality. The base-case mortality parameter came from the tier 1 instrumental-variable analysis of Basu et al[11] (county-level instrumental variable using National Health Service Corps loan-forgiveness purchasing power; first-stage F=25.6). We summed cause-specific estimates (cardiovascular, cancer, respiratory, infectious disease, and substance use/injury) to obtain a composite all-cause effect of −1.20 deaths per 100,000 population per additional physician per 100,000 population (95% CI, −1.78 to −0.63). This specification was chosen a priori for its causal rigor and direct applicability to the simulation metric.

Preventable hospitalizations. The preventable hospitalization parameter was harmonized from the composite PQI-90 estimate in McMichael[14] using staggered difference-in-differences based

on state nurse practitioner full-practice-authority adoption. The harmonized effect was −0.608 admissions per 100,000 population per additional provider per 100,000 population (95% CI, −0.945 to −0.273; eAppendix A).

Sensitivity Analyses

Two sensitivity analyses were prespecified. First, the physician-equivalence factor (e) was varied from 25% to 100% of full physician effectiveness across all 3 adoption scenarios (Table 3). Second, the moderate adoption scenario was rerun using 5 alternative mortality parameters from the systematic review, spanning conservative observational panel estimates to upper-bound instrumental-variable specifications (eAppendix C). To assess geographic heterogeneity, outcomes were decomposed by HPSA designation status, urban/rural classification, and Census region (Table 2).

Statistical Analysis

Confidence intervals for projected outcomes were obtained by analytically propagating the published 95% CIs of the causal parameters through the simulation while holding all other inputs (county supply, q^*, τ, and equivalence factor) at point estimates. Because outcomes are linear in the causal coefficients and adoption rates, analytic propagation yields results identical to Monte Carlo simulation with infinite draws and avoids unnecessary simulation noise. A consolidated table of all simulation parameters, base-case values, uncertainty ranges, and sources appears in eAppendix C. This decision analytical model followed the relevant non-cost aspects of the Consolidated Health Economic Evaluation Reporting Standards (CHEERS 2022) reporting guideline (eAppendix C provides model structure, assumptions, input parameters, and uncertainty analysis). Analyses were conducted using Stata, version 18 (StataCorp LLC).

Results

Primary Care Supply and the Treated Population

The analytic sample comprised 3,106 counties. The supply threshold q* was set at 66.7 physicians per 100,000 population (1:1,500 ratio). A total of 2,303 counties (74.1%) fell below q*, encompassing 141.86 million residents. The county supply distribution was right-skewed (Figure 1): the median county had 48.3 physicians per 100,000 population, and the bottom quartile had fewer than 30, falling below the HRSA HPSA designation threshold. Figure 1 displays the distribution and the q* threshold alongside 2 reference benchmarks: the HRSA HPSA primary care designation cutoff (28.6 per 100,000)[15] and the Council on Graduate Medical Education (COGME)–recommended generalist supply (80 per 100,000).[16]

Base-Case Results

Under the moderate adoption scenario (τ=0.31), deployment of an AI PCP yielded 8,922 effective physician-equivalents added across 2,303 treated counties, reaching an estimated 43.98 million patients. Deployment was associated with an estimated 10,707 deaths averted per year (95% CI, 5,621-15,882) and 5,422 preventable hospitalizations averted per year (95% CI, 2,417-8,432) (Table 1). Under the conservative scenario (τ=0.19), estimated deaths averted were 6,562 (95% CI, 3,445-9,734) and preventable hospitalizations averted were 3,323 (95% CI, 1,481-5,168). Under the upper scenario (τ=0.47), estimated deaths averted were 16,233 (95% CI, 8,522-24,079) and preventable hospitalizations averted were 8,221 (95% CI, 3,665-12,784).

Geographic Distribution

The geographic concentration of effective physician-equivalents added tracked the underlying distribution of supply gaps (Figure 2). Under the base case (moderate adoption), the South accounted for 54.0% of estimated deaths averted (5,778 of 10,707) and an identical share of preventable hospitalizations averted (Table 2). This concentration reflects the South's disproportionate representation among shortage counties (1,171 of 2,303 treated counties [50.8%]) and its larger at-risk population (67.26 of 141.86 million treated residents [47.4%]).

Nonmetropolitan counties accounted for 34.9% of estimated deaths averted despite housing only 25.4% of the treated population, consistent with the deeper average supply shortfalls in rural areas. The Northeast, with the fewest shortage counties (111) and the smallest treated population (17.19 million), accounted for just 7.4% of simulated benefit. HPSA-designated counties accounted for 30.5% of estimated outcomes.

Sensitivity to Equivalence and Parameter Assumptions

Table 3 displays projected annual deaths and preventable hospitalizations averted across 4 physician-equivalence factors and 3 adoption scenarios. Even under the most skeptical reasonable position—conservative adoption (τ=0.19) combined with one-quarter-physician equivalence (e=0.25)—the projected annual mortality reduction was 1,641 deaths and the projected reduction in preventable hospitalizations was 831 admissions. At moderate adoption, halving the equivalence assumption (e=50%) yielded 5,353 deaths and 2,711 preventable hospitalizations averted annually. Results were also directionally consistent across 5 alternative mortality parameters drawn from the systematic review (eAppendix C), ranging from approximately 1,200 to 10,707 deaths averted under moderate adoption.

Discussion

In this county-level simulation, national deployment of an autonomous AI PCP in shortage counties—at adoption rates consistent with current telehealth utilization and assuming physician-equivalent clinical effectiveness—was projected to avert approximately 10,700 deaths and 5,400 preventable hospitalizations annually. These estimates remained policy-relevant even under substantially more conservative assumptions about AI effectiveness (Table 3) and were directionally consistent across multiple alternative causal parameters (eAppendix C). Benefits were concentrated in regions with the deepest existing shortfalls: the South accounted for more than half of projected deaths averted, and nonmetropolitan counties were

overrepresented relative to their share of the treated population. In magnitude, the projected mortality reduction under moderate adoption is roughly twice the annual reduction linked to Affordable Care Act Medicaid expansion among low-income older adults[17] and comparable in scale to modeled benefits from widespread glucagon-like peptide-1 (GLP-1) receptor agonist use.[18]

The physician-equivalence assumption is central to these projections and merits careful scrutiny. A structured review of head-to-head studies published between January 2023 and April 2026 comparing frontier LLMs against practicing physicians on primary care–relevant reasoning tasks is reported in eAppendix B; the central tendency places frontier LLMs at or above physician performance on diagnostic and management reasoning. Randomized double-blind crossover OSCE evaluations rated AI systems superior to practicing primary care physicians on the large majority of specialist- and patient-actor–evaluated axes,[7] and the first prospective real-world deployment of an autonomous LLM-based system in an ambulatory primary care clinic found no significant differences from primary care physicians in differential diagnosis quality, management appropriateness, or safety.[19] Although real-world evidence on downstream outcomes such as mortality remains absent, adopting physician-equivalent effectiveness as the base-case modeling convention is consistent with current evidence on intermediate clinical outcomes.

The populations projected to benefit most are those with the highest baseline barriers to digital health engagement. Rural communities face disproportionate barriers—such as broadband access gaps and lower digital health literacy—that current deployment models only partially address.[20,21] Without deliberate implementation strategies addressing these barriers, the projected geographic concentration of benefit, which reflects the distribution of need, could in practice be muted by the distribution of access.

Health insurance coverage and physician supply are spatially correlated. Southern and rural counties facing the deepest supply gaps also face the highest uninsurance rates and the most restrictive Medicaid programs, raising the question of whether the simulation conflates the two—whether outcome differences attributed to supply are partly attributable to coverage. Both parameters, however, are estimated in real-world US samples in which coverage already operates as a limiting factor. The per-physician supply–mortality estimate is identified from variation in physician supply in a population whose ability to access those physicians is determined by the coverage distribution as it actually exists; coverage gaps therefore attenuate the realized per-physician estimate that the underlying causal studies recover. The adoption rate is similarly calibrated to observed telehealth utilization in a survey population carrying the existing coverage mix. The residual coverage-related uncertainty therefore operates through a feature of the AI PCP we have not specified—its cost-sharing structure—rather than through a parameter to be estimated. A platform deployed free at the point of use would reach uninsured persons at higher rates than the marginal physician in the supply–outcome literature, whereas a platform requiring commercial insurance reimbursement would attenuate projected benefit in precisely the Southern and rural geographies where the model concentrates it.

Cost Considerations

A bottom-up estimate of deployment cost can be triangulated from 3 sources of contemporary pricing data. Frontier LLM inference, at published 2026 application programming interface (API) rates,[22] costs approximately $1 to $3 per patient per year for a typical primary care utilization profile. AI-first ambulatory deployments operating at scale—such as Doctronic, which has completed more than 25 million consultations—deliver AI consultation at zero marginal cost to the patient, with human telehealth escalation capped at $39 per visit.[23] Comprehensive virtual primary care via established telehealth platforms is delivered at approximately $130 per member per year.[24] Triangulating these benchmarks places the all-in annual deployment cost at

$100 to $150 per patient (plausible range, $50-$300), corresponding to approximately $4 billion to $7 billion annually for the 44 million patients reached under moderate adoption. Against the projected 10,707 averted deaths annually, this corresponds to an implied cost of approximately $470,000 per death averted in the central case (range, $200,000 to $1.2 million across cost scenarios). This calculation is dominated by the fact that physician-equivalent cognitive labor, once developed, has near-zero marginal cost; deployment cost scales with platform delivery rather than with consultation volume.

These projections arrive as the regulatory framework for autonomous clinical AI is actively being developed. The US Food and Drug Administration Software as a Medical Device framework was designed for static, task-bounded tools and is widely recognized as inadequate for adaptive, general-purpose AI systems capable of open-ended clinical reasoning.[25] A licensure-based framework for autonomous clinical AI—requiring standardized competency examination, supervised deployment, and time-limited federal certification—has recently been proposed.[26] The projections reported here quantify the scale of benefit such a framework would need to govern responsibly and the scale of harm a poorly calibrated framework would risk allowing through.

Limitations

This study has several limitations. First, the simulation assumes physician-equivalent clinical effectiveness; no randomized trials of autonomous AI primary care with long-term patient outcomes currently exist, and head-to-head evidence remains stronger for single-condition reasoning than for the complex, multimorbid longitudinal care that predominates in shortage areas. Second, telehealth-based adoption rates may not fully translate to AI-specific engagement, particularly in populations with limited digital infrastructure or trust. Third, the supply threshold q* was set at 66.7 physicians per 100,000 population (1:1,500 ratio); alternative thresholds spanning the HRSA HPSA designation cutoff (28.6) through the COGME-

recommended generalist supply (80) and an above-adequacy benchmark (100) are reported in eAppendix D and produce direction-consistent estimates that scale with the supply gap. Fourth, the simulation abstracts from general equilibrium effects, including substitution away from existing primary care, changes in referral patterns, and induced demand. Fifth, the model applies the causal supply–outcome relationship beyond the exact range of supply variation observed in the source studies, so point estimates should be interpreted as upper-bound projections under local linearity. Sixth, the simulation does not model downstream specialist capacity; the reported estimates therefore assume specialist capacity is sufficient to absorb induced referral volume, and to the extent it is not, realized mortality and preventable-hospitalization benefits will be attenuated, particularly in rural counties where specialist density is lowest.

Conclusions

National deployment of autonomous AI primary care as a supply-side complement in shortage counties has the potential to avert meaningful population health harms at a scale relevant to federal policy. Realizing this potential will require rigorous prospective validation in multimorbid populations, thoughtful payment and regulatory design, and implementation strategies that actively address digital access and trust barriers in the communities that stand to benefit most.

Article Information

Author Contributions: Dr Bergman had full access to all of the data in the study and takes responsibility for the integrity of the data and the accuracy of the data analysis. Concept and design, acquisition, analysis, or interpretation of data, drafting of the manuscript, critical revision of the manuscript for important intellectual content, statistical analysis, administrative, technical, or material support, and supervision: Bergman.

Conflict of Interest Disclosures: None reported.

Funding/Support: None.

Role of the Funder/Sponsor: Not applicable.

Data Sharing Statement: All data used in this study are publicly available from the sources listed in the Methods section, including the Area Health Resources File (HRSA), CDC WONDER, the HRSA HPSA Shortage Designation file, and USDA Rural-Urban Continuum Codes. Simulation code and analytic files will be made available upon publication upon reasonable request to the corresponding author.

Additional Information: This article includes eAppendix A (systematic review of causal estimates of the primary care supply–outcome relationship), eAppendix B (structured review of head-to-head AI–physician studies), eAppendix C (consolidated simulation input parameters), and eAppendix D (sensitivity analyses for alternative thresholds and causal parameters), available as Supplement.

References

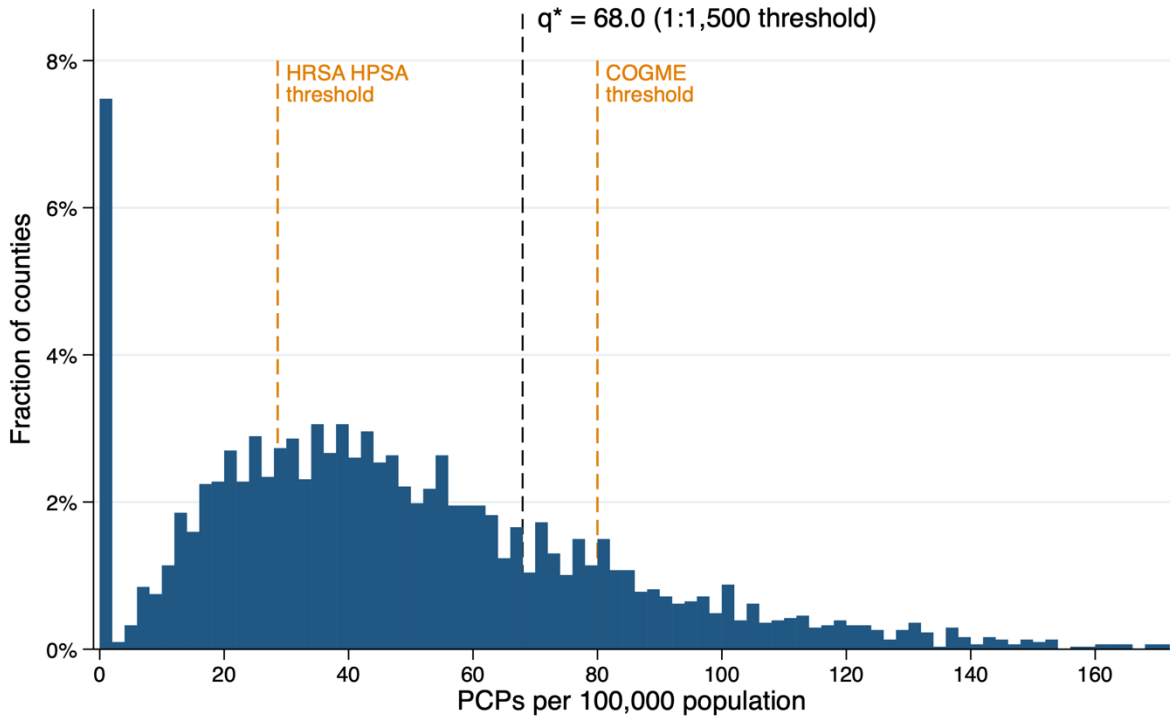
1. Designated Health Professional Shortage Areas statistics, 2025. Health Resources and Services Administration, HRSA Data Warehouse. Accessed April 23, 2026. https://data.hrsa.gov/topics/health-workforce/shortage-areas
2. Health Workforce Projections: 2023-2038. Health Resources and Services Administration, National Center for Health Workforce Analysis; 2025.
3. Horstman C, Shah A. The State of Rural Primary Care in the United States. Commonwealth Fund; 2025.
4. Dwyer-Lindgren L, Bertozzi-Villa A, Stubbs RW, et al. US county-level trends in mortality rates for major causes of death, 1980-2014. JAMA. 2016;316(22):2385-2401.
5. Singhal K, Azizi S, Tu T, et al. Large language models encode clinical knowledge. Nature. 2023;620(7972):172-180.
6. Kanjee Z, Crowe B, Rodman A. Accuracy of a generative artificial intelligence model in a complex diagnostic challenge. JAMA. 2023;330(1):78-80.
7. Tu T, Schaekermann M, Palepu A, et al. Towards conversational diagnostic artificial intelligence. Nature. 2025;642(8067):442-450.
8. Area Health Resources File (AHRF), 2025 release. Health Resources and Services Administration; 2025.
9. Multiple Cause of Death Data, 1999-2024. Centers for Disease Control and Prevention, CDC WONDER. Accessed 2025. https://wonder.cdc.gov
10. Rural-Urban Continuum Codes (RUCC), 2023. US Department of Agriculture Economic Research Service; 2023.

11. Basu S, Berkowitz SA, Phillips RL, Bitton A, Landon BE, Phillips RS. Association of primary care physician supply with population mortality in the United States, 2005-2015. JAMA Intern Med. 2019;179(4):506-514.
12. Telehealth Utilization Among US Adults: Nationally Representative Survey (SSRS, March-August 2023). Commonwealth Fund; 2025.
13. Sun R, Blackburn B, Aledhaim A, et al. Trends in telehealth utilization among Medicare fee-for-service beneficiaries during the COVID-19 pandemic. medRxiv. Preprint posted online 2025. doi:10.1101/2025.03.05.25323449
14. McMichael BJ. Nurse practitioner full practice authority and preventable hospitalizations: evidence from state-level policy changes. J Health Econ. Published online ahead of print 2025.
15. Designation of Health Professional(s) Shortage Areas. 42 CFR §5, Appendix A. US Department of Health and Human Services. Accessed May 5, 2026. https://www.ecfr.gov/current/title-42/chapter-I/subchapter-A/part-5
16. Patient Care Physician Supply and Requirements: Testing COGME Recommendations. Eighth Report. Council on Graduate Medical Education, Health Resources and Services Administration; 1996.
17. Miller S, Johnson N, Wherry LR. Medicaid and mortality: new evidence from linked survey and administrative data. Q J Econ. 2021;136(3):1783-1829.
18. Walsh T, Guzman G. GLP-1 therapies and mortality risk: implications for life insurers. Munich Re Life US. February 12, 2026. Accessed April 27, 2026. https://www.munichre.com/us-life/en/insights/clinical-knowledge/glp-1-therapies-and-mortality-risk-implications-for-life-insurers.html

19. Schaekermann M, Tu T, Palepu A, et al. A prospective clinical feasibility study of a conversational diagnostic AI in an ambulatory primary care clinic. arXiv. Preprint posted online 2026. arXiv:2603.08448
20. Nouri S, Khoong EC, Lyles CR, Karliner L. Addressing equity in telemedicine for chronic disease management during the Covid-19 pandemic. NEJM Catal Innov Care Deliv. 2020;1(3).
21. Luo J, Tong L, Crotty BH, et al. Telemedicine adoption during the COVID-19 pandemic: gaps and inequalities. JAMA Intern Med. 2022;182(8):835-844.
22. Claude API pricing documentation. Anthropic. Accessed May 5, 2026. https://docs.claude.com/en/docs/about-claude/pricing
23. Pricing and resource center. Doctronic. Accessed May 5, 2026. https://www.doctronic.ai
24. Annual Report on Form 10-K: Subscription Access Fee Disclosures. Teladoc Health Inc; 2024. Accessed May 5, 2026. https://ir.teladochealth.com
25. Software as a Medical Device (SaMD). US Food and Drug Administration. Accessed April 23, 2026. https://www.fda.gov/medical-devices/digital-health-center-excellence/software-medical-device-samd
26. Bergman A, Wachter RM, Emanuel EJ. A licensure framework for autonomous clinical AI. JAMA. Published online April 28, 2026.

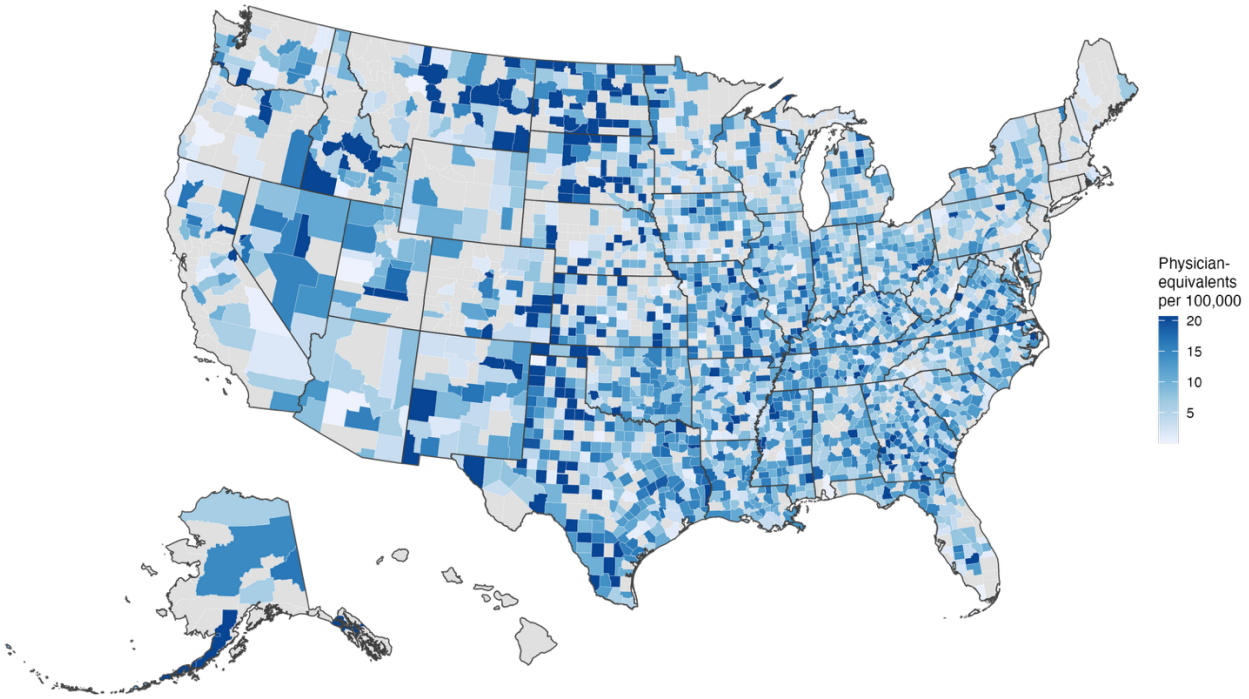
Exhibits

Figure 1. Distribution of County-Level Primary Care Physician Supply and Simulation Deployment Threshold.



The histogram displays the distribution of county-level primary care physician supply (physicians per 100,000 population) for the 3,106-county analytic sample. The vertical dashed line marks the deployment threshold q* (66.7 per 100,000 population), equivalent to a 1:1,500 population-to-physician ratio. US territories are excluded.

Figure 2. County-Level Effective Physician-Equivalents Added per 100,000 Population Under Moderate AI Primary Care Adoption.



The choropleth map shows effective physician-equivalents added per 100,000 population under the moderate adoption scenario (τ=0.31), computed as τ × (q* − observed county primary care physician supply), with q*=66.7 per 100,000 population. Counties at or above q* are shaded gray and contribute zero.

Table 1. Simulated Annual Outcomes Under 3 AI Primary Care Adoption Scenarios

	Conservative (τ = 0.19)	Moderate (τ = 0.31)	Upper (τ = 0.47)
Population in treated counties, millions	141.86	141.86	141.86
Population using AI PCP, millions[a]	26.95	43.98	66.67
Effective physician-equivalents added	5,469	8,922	13,528
Deaths averted per year (95% CI)[b]	6,562 (3,445-9,734)	10,707 (5,621-15,882)	16,233 (8,522-24,079)
PQI-90 admissions averted per year (95% CI)[b]	3,323 (1,481-5,168)	5,422 (2,417-8,432)	8,221 (3,665-12,784)

Abbreviations: AI PCP, autonomous artificial intelligence primary care physician; CI, confidence interval; PQI-90, Prevention Quality Indicator composite (Agency for Healthcare Research and Quality).
All 2,303 counties with primary care physician supply below q* (66.7 per 100,000 population; 1:1,500 ratio) are included in each scenario. Scenarios differ in the fraction of eligible patients who adopt the AI PCP (τ).
[a] Population using the AI PCP equals the population in treated counties multiplied by τ.
[b] Mortality estimates were derived from Basu et al[11] using a county-level instrumental variable design; preventable hospitalization estimates were derived from McMichael[14] using a difference-in-differences design with state-level nurse practitioner full-practice-authority adoption as the source of quasi-experimental variation.

Table 2. Geographic Distribution of Projected Annual Outcomes Under Moderate AI Primary Care Adoption (Base Case)

Subgroup	Counties, No.	Population, millions	Physician-equivalents added	Deaths averted	PQI-90 admissions averted	% of total
HPSA designation status						
Non-HPSA counties	1,418	105.54	6,205	7,447	3,771	69.6
HPSA-designated counties	885	36.32	2,717	3,260	1,651	30.5
Urban/rural classification						
Metropolitan (RUCC 1-3)	755	105.84	5,804	6,965	3,527	65.1
Nonmetropolitan (RUCC 4-9)	1,539	36.01	3,116	3,740	1,894	34.9
Census region						
Northeast	111	17.19	657	788	399	7.4
Midwest	753	28.93	2,132	2,559	1,296	23.9
South	1,171	67.26	4,815	5,778	2,926	54.0
West	268	28.48	1,318	1,582	801	14.8
All treated counties	**2,303**	**141.86**	**8,922**	**10,707**	**5,422**	**100.0**

Abbreviations: AI, artificial intelligence; HPSA, Health Professional Shortage Area; PQI-90, Prevention Quality Indicator composite; RUCC, Rural-Urban Continuum Code.

Base-case scenario: moderate adoption (τ=0.31). Counties were classified as HPSA-designated if they appeared in the HRSA HPSA Shortage Designation file with a score of 1 or greater for primary care. Urban/rural classification was based on USDA Rural-Urban Continuum Codes (metropolitan, RUCC 1-3; nonmetropolitan, RUCC 4-9). Census region assignments follow US Census Bureau definitions. The percentage column reflects each subgroup's share of the simulation total for deaths averted (identical to its share of preventable hospitalizations averted under the linear model).

Table 3. Projected Annual Deaths and Preventable Hospitalizations Averted by Physician-Equivalence Factor and Adoption Scenario

Equivalence factor (e)	Conservative (τ = 0.19)	Moderate (τ = 0.31)	Upper (τ = 0.47)
Panel A. Deaths averted per year			
100% (base case)	6,562	10,707	16,233
75%	4,922	8,030	12,175
50%	3,281	5,353	8,117
25%	1,641	2,677	4,058
Panel B. PQI-90 admissions averted per year			
100% (base case)	3,323	5,422	8,221
75%	2,492	4,067	6,166
50%	1,662	2,711	4,110
25%	831	1,356	2,055

Abbreviations: AI, artificial intelligence; PQI-90, Prevention Quality Indicator composite.
The equivalence factor e represents the share of a primary care physician's outcome estimate that the AI primary care physician is assumed to deliver. The base case (e=100%) assumes the AI primary care physician is as effective as a physician; e=50% assumes half as effective. Adoption scenarios (τ) reflect the share of eligible patients who use the AI primary care physician, calibrated to observed rural primary care telehealth utilization (conservative), nonrural primary care telehealth utilization (moderate), and the peak primary care telehealth share among Medicare fee-for-service beneficiaries during April 2020 (upper). Each cell shows projected national totals across the 2,303 counties below the supply threshold q* (66.7 per 100,000 population; 1:1,500 ratio). Mortality estimates were derived from Basu et al[11]; preventable hospitalization estimates were derived from McMichael[14].