



WHITE PAPER · SEPTEMBER 2026

Pro-Worker AI in Health

Who benefits when machines
enter the largest workforce in
America?

Blythe Adamson, PhD, MPH

INFECTIOUS ECONOMICS



Photo: Jakub Halun, CC BY 4.0, via Wikimedia Commons



Pro-Worker AI in Health: Who Benefits When Machines Enter the Largest Workforce in America?

Blythe Adamson, PhD, MPH

Founder, Infectious Economics, LLC, New York, NY

blythe@infectious-economics.com

Suggested citation. Adamson B. Pro-Worker AI in Health: Who Benefits When Machines Enter the Largest Workforce in America? Infectious Economics; September 2026. Available at: <https://blytheadamson.com/pro-worker-ai-in-health>. doi:10.2139/ssrn.7466599

About this report. This white paper applies the Acemoglu, Autor, and Johnson (2026) pro-worker AI framework to a structured environmental scan of 132 AI-enabled health technologies deployed in the United States. It was prepared for policymakers, health system leaders, investors, and the workforce whose jobs the technologies touch. The full technology-level dataset is reproduced in Appendix B and is available for download, under a CC BY 4.0 license, at blytheadamson.com/pro-worker-ai-in-health.

Version. Version 1, September 12, 2026.

Funding and independence. This work received no external funding and was supported by Infectious Economics, LLC. Disclosures appear at the end of the report.

Photography. Cover: Jakub Hafun (CC BY 4.0). Section images: Paulo H. Carvalho, Agência Brasília (CC BY 2.0); U.S. Navy (public domain); James Gathany, CDC Public Health Image Library (public domain); Diane A. Reid, National Cancer Institute (public domain); DIAC images (CC BY 2.0). All via Wikimedia Commons. Photographs are illustrative and do not depict any technology in the scan.

© 2026 Infectious Economics, LLC. This report may be shared and quoted with attribution.

CONTENTS

A guide to the argument

Key findings	4
The report in brief	5
The argument on one page	7
Who benefits, and who decides	8
01 Why health AI is also a labor-market question	9
02 What makes a technology pro-worker?	12
03 What the product landscape shows	16
04 Who captures the productivity gains?	28
05 Five levers for beneficial deployment	38
06 Limits of the evidence, and conclusions	45
About the author and disclosures	48
References	49
Appendix A Supplementary methods	51
Appendix B Technology-level inventory	53

Reading guide: the key findings, the report in brief, and the two one-page summaries give the argument in ten minutes. Sections 01 to 06 carry the evidence and the reasoning. Appendix B reproduces every technology in the scan so the classifications can be audited or re-coded.

KEY FINDINGS

- 01** **The pro-worker character of health AI is a design and policy variable, not a technological given.** The same capabilities can script and monitor a worker, augment her judgment, or equip her to do work no one was doing.
- 02** **Within a scan of 132 deployed technologies, administrative AI is almost entirely automation.** 39 of 40 provider- and payer-administration products (97.5%) were classified as automating existing tasks.
- 03** **Health-producing domains look different.** 61 of 92 clinical, life sciences, and public health products (66%) were classified as creating new tasks or extending expertise to new workers and settings.
- 04** **The friction-cost paradox.** The workers most exposed to automation do work that exists because of billing and insurance complexity. Reducing it is efficient for the system and costly for a predominantly female clerical workforce.
- 05** **Who captures the gains depends on market structure.** Payment incentives, employer concentration, insurance pass-through, and scope-of-practice rules decide whether productivity gains reach workers and patients.
- 06** **Five policy levers can change the trajectory,** each with a responsible actor and an evaluation criterion. Support should follow demonstrated benefit, not novelty.

EXECUTIVE SUMMARY

Health care is the largest employer in America. AI is entering it unevenly.

The US health sector accounts for 18.0% of gross domestic product and roughly one in eight jobs. The artificial intelligence products now being deployed across it will change what its workers do, but the labor consequences of those products have never been inventoried. This report is a first, descriptive answer.

132

AI-enabled health technologies deployed or FDA-authorized in the US, classified under a pre-specified protocol

97.5%

of administrative and payer technologies classified as automating existing tasks (39 of 40)

66%

of clinical, life sciences, and public health technologies classified as new task-creating or expertise-leveling (61 of 92)

92M

Americans living in a primary care shortage area at the end of 2025, where expertise-leveling tools could expand access

I applied the Acemoglu, Autor, and Johnson (2026) five-category pro-worker framework to a structured environmental scan of 132 commercially deployed or FDA-authorized AI-enabled health technologies identified in March 2026 across five domains: clinical care, provider administration, payer administration, life sciences, and public health. Each technology was classified by the anticipated direction of its effect on labor productivity, the value of human expertise, and labor's share of income.

Within the scan, 41.7% of technologies were classified as automating existing tasks, 29.5% as creating new tasks, 16.7% as expertise-leveling, 10.6% as labor-augmenting, and 1.5% as capital-augmenting. The classifications diverge sharply by domain. Administrative products are almost all automation. Clinical, life sciences, and public health products mostly create new tasks or extend expertise to new workers and settings.

These are classifications of anticipated consequences, not measurements of employment, wages, or income shares. I offer an economic interpretation as a set of testable hypotheses: administrative tasks are rule-bound and algorithmically reducible, whereas clinical and scientific judgment sustains demand for human expertise, and the distribution of any productivity gains depends on payment incentives, provider and employer market power, and workforce regulation.

I describe a friction-cost paradox, in which the workers most exposed to automation perform work that exists largely because of insurance and billing complexity, and I propose five policy levers with responsible actors and evaluation criteria. The scan is limited to commercially oriented products and cannot support conclusions about net employment effects or about the level of investment in public health AI.

EXHIBIT 1

The task mix differs sharply across domains: administrative AI is almost entirely automation, while health-producing domains mostly create new tasks or extend expertise.

Count and within-domain share of 132 AI-enabled US health technologies by AAJ pro-worker category, March 2026. Darker cells indicate a larger share of the domain.

	Automation	New task-creating	Expertise-leveling	Labor-augmenting	Capital-augmenting
Clinical care n = 56	13 23.2%	14 25.0%	15 26.8%	12 21.4%	2 3.6%
Provider administration n = 26	25 96.2%	0 0.0%	0 0.0%	1 3.8%	0 0.0%
Payer administration n = 14	14 100.0%	0 0.0%	0 0.0%	0 0.0%	0 0.0%
Life sciences n = 21	0 0.0%	20 95.2%	0 0.0%	1 4.8%	0 0.0%
Public health n = 15	3 20.0%	5 33.3%	7 46.7%	0 0.0%	0 0.0%

Source: Infectious Economics environmental scan, March 22-28, 2026 (Appendix B). Public health includes population health and workforce products.

The same natural language processing, predictive analytics, and computer vision can be configured to script and surveil a worker, to augment her judgment, or to equip her with entirely new diagnostic capabilities. What varies is the role assigned to human expertise.

VISUAL ABSTRACT

The argument on one page

Pro-Worker AI in Health

Who benefits when machines enter the largest workforce in America?

VISUAL ABSTRACT

Adamson B. Infectious Economics, 2026

1. What was studied

132 AI-enabled health technologies deployed or FDA-authorized in the US

Structured environmental scan, March 22 to 28, 2026

Five domains

Clinical care	n = 56
Provider administration	n = 26
Payer administration	n = 14
Life sciences	n = 21
Public health	n = 15

Each classified with the Acemoglu, Autor and Johnson (2026) pro-worker taxonomy by its anticipated effect on productivity, the value of expertise, and labor's share:

New task-creating	pro-worker
Expertise-leveling	ambiguous
Labor-augmenting	ambiguous
Capital-augmenting	ambiguous
Automation	not pro-worker

2. What the scan shows

Administrative layer (provider + payer) n = 40

39 automation, 1 labor-augmenting

97.5% classified as automation of existing tasks (39 of 40)

clinical care, life sciences, public health

Health-producing domains n = 92

39 new task-creating, 22 expertise-leveling, 13 labor-augmenting, 2 capital-augmenting, 16 automation

66% classified as new task-creating or expertise-leveling (61 of 92)

Robust to re-coding ambient scribes as labor-augmenting and life-science tools as labor-augmenting; 39 of 40 administrative tools remain automation either way.

Classifications are anticipated consequences, not observed employment or wage outcomes.

3. Who captures the gains

The pattern tracks market structure, not the technology. Whether gains reach workers depends on:

- **Payment incentives**
fee-for-service vs value-based
- **Labor market concentration**
employer power over wages
- **Insurance pass-through**
savings to premiums or margin
- **Workforce regulation**
scope of practice and licensure

Five policy levers

- 1 Align coding, coverage and payment with benefit
- 2 Report administrative AI savings and allocation
- 3 Supervised scope-of-practice pilots
- 4 Public demand-pull for public health AI
- 5 Transition support and worker voice

The friction-cost paradox

The workers most exposed to automation perform work that exists because of billing and insurance complexity. Reducing that work is efficient for the system and costly for a predominantly female clerical workforce (women are 92% of medical records specialists). Welfare gains are conditional on preserving valuable administrative functions and managing the transition.

Framework: Acemoglu D, Autor D, Johnson S. Building Pro-Worker Artificial Intelligence. NBER Working Paper 34854, 2026. Data: Supplementary Table S1.

blytheadamson.com

Counts from Appendix B. Women's share of medical records specialists from the Bureau of Labor Statistics Current Population Survey, 2025 annual averages.

CONCEPTUAL FRAMEWORK

Who benefits from health AI, and who decides

Deployment choices determine what a technology does to tasks.
Institutions determine who captures the gains.

Administration

Automation of existing tasks. Efficiency gains can concentrate worker risks in a predominantly female clerical workforce.

39 of 40 automation

Clinical care and life sciences

Augmentation, new tasks, and extended expertise. Effects vary by occupation and by setting.

56 of 77 new task or leveling

Public health

Surveillance and community capabilities. Procurement requires sustained operating support to be useful.

5 of 15 new task-creating

Institutions shape who captures the gains

Payment incentives · Labor market concentration · Insurance pass-through · Workforce regulation · Trust and accountability

OUTCOMES TO EVALUATE

Workers

Value of expertise, earnings, employment, and job quality in the affected occupations.

Patients and communities

Access, quality, affordability, and time to detection of population threats.

Capital owners

Margins and returns, including savings retained rather than passed through.

Conceptual pathways, not causal estimates. The scan classifies anticipated effects; measured employment and wage outcomes require further study. Clinical care and life sciences pooled: 14 + 20 new task-creating and 15 + 0 expertise-leveling of 77 products.

A close-up portrait of a woman with short, wavy brown hair, smiling warmly. She is wearing blue medical scrubs and a black stethoscope around her neck. A name tag on her chest reads "Gayle" with a small logo. The background is a blurred clinical setting with white cabinets and medical equipment.

01

The largest workforce in America

Health care employs one in eight American workers. Changes in what those workers do have macroeconomic consequences, and the AI products now entering the sector have never been inventoried for their labor effects.

01 • INTRODUCTION

Why the question matters now

National health expenditure grew 7.2% to \$5.3 trillion in 2024, or \$15,474 per person, and accounted for 18.0% of gross domestic product (CMS 2025). The Centers for Medicare and Medicaid Services (CMS) Office of the Actuary projects that health spending will reach 20.3% of GDP by 2033 (Keehan et al. 2025). Health care and social assistance is the largest employment sector in the United States, and health occupations alone account for roughly one in eight jobs (BLS 2023). Changes in labor demand within this sector therefore have macroeconomic consequences. A growing literature examines how artificial intelligence (AI) affects labor markets in general (Acemoglu and Restrepo 2018, 2019; Autor 2024), and a large clinical literature evaluates the accuracy and safety of individual AI tools. I am not aware, however, of a cross-domain inventory of deployed US health AI products that classifies each product by its anticipated consequences for the workers who perform the affected tasks. That is the narrow gap this paper addresses.

I apply the Acemoglu, Autor, and Johnson (2026) five-category framework for pro-worker AI (hereafter AAJ) to a structured environmental scan of 132 commercially deployed or FDA-authorized AI-enabled health technologies spanning five domains: clinical care, provider administration, payer administration, life sciences, and public health. For each technology I assessed the anticipated direction of its effect on labor productivity, on the value of human expertise, and on labor's share of income, and classified it into one of the five AAJ categories. I then interpret the pattern of classifications through the institutional features of US health care markets that health economists have long identified as shaping how productivity gains are generated and captured: information asymmetry (Arrow 1963), provider consolidation and employer market power (Gaynor and Town 2011; Brot-Goldberg et al. 2024), the entanglement of health insurance with employment (Gruber 2000), provider payment incentives (Chernew 2020), and workforce regulation (Nicholson and Propper 2011).

Two boundaries on the contribution should be stated at the outset. First, the scan measures my classification of selected products against a pre-specified decision rule. It does not measure adoption intensity, workers displaced, jobs created, compensation, or the effect of alternative payment arrangements. Throughout the paper I distinguish three kinds of statements: observed product characteristics, theoretical predictions drawn from the AAJ framework, and policy hypotheses. Second, the domain differences I report are patterns within this scan. The economic explanations I propose for them are hypotheses that require evaluation with employment and earnings data, not conclusions established here.

The central pattern is that the classifications vary systematically by domain. Provider and payer administration, which I call the friction-cost layer because much of the labor it employs exists to manage billing, claims, and insurance complexity, is almost entirely classified as automating existing tasks. Clinical care, life sciences, and public health, where labor contributes directly to patient outcomes or population health, are disproportionately classified as creating new tasks or extending expertise to new workers and settings. If this pattern reflects the underlying economics of each domain, as I argue it plausibly does, then the pro-worker character of health AI is shaped less by the technology than by the market structure, payment incentives, and regulatory context into which it is introduced, and it is therefore amenable to policy.

THREE KINDS OF STATEMENT IN THIS REPORT

Observed product characteristics

What a technology does, where it is deployed, and its regulatory status, as recorded in the dataset.

Theoretical predictions

The anticipated direction of effect on productivity, expertise, and labor's share, drawn from the AAJ framework.

Policy hypotheses

Claims about market structure, payment incentives, and who captures the gains, to be tested with employment and earnings data.



02

What makes a health AI technology pro-worker?

A precise definition, and three illustrative deployments for community health workers that show how wide the range of labor consequences can be for a single occupation.

02 • DEFINITION AND SCOPE

Pro-worker has a precise meaning in this report



The AAJ framework classifies technologies by their effect on the tasks workers perform and on the value of the expertise those tasks require. A technology is labor-augmenting when it makes workers more productive at tasks they already perform, capital-augmenting when it makes machines more productive at their tasks, automating when it substitutes machines for tasks that workers previously performed, new task-creating when it generates demand for human expertise in tasks that did not previously exist, and expertise-leveling when it enables less-expert workers to perform tasks that previously required greater expertise. In the AAJ account only new task-creating technologies are unambiguously pro-worker, because only they expand the set of tasks in which human expertise commands a return. Automation is not pro-worker in this sense because it commodifies expertise. The other three categories are ambiguous: their net effect depends on demand responsiveness, bargaining power, and the pace at which incumbents' scarcity rents erode.

Because the term is used loosely in policy debate, I state how it is used here. The primary meaning of pro-worker in this paper is the AAJ meaning: a technology is pro-worker to the extent that it raises the value of human expertise and expands labor's share of the income generated by the affected activity. The workers whose outcomes count are those in the occupations that perform the affected tasks, identified by Standard Occupational Classification (SOC) code in the scan dataset. The horizon is the near to medium term, roughly the next decade, over which current products are deployed at scale. Three distinctions follow. Task displacement is not job loss: an occupation can lose tasks and retain or gain employment if demand for its remaining tasks grows. Wage levels are not labor's share: wages can rise for some workers while the share of income accruing to labor falls. And worker outcomes are not total welfare: a technology can be good for patients and payers while being bad for a specific workforce. Where the paper discusses burnout, access, or welfare, it says so explicitly rather than folding those outcomes into the term pro-worker.

EXHIBIT 2

Only new task-creating technologies are unambiguously pro-worker in the AAJ framework.

The five categories, what they do to tasks, and their anticipated effect on the value of human expertise and labor's share of income

 New task-creating Generates demand for human expertise in tasks that did not previously exist in the occupation or setting. UNAMBIGUOUSLY PRO-WORKER	 Expertise-leveling Lets less-expert workers, or workers in new settings, perform tasks that required greater expertise. AMBIGUOUS	 Labor-augmenting Makes workers more productive at tasks they already perform; the tasks stay with the worker. AMBIGUOUS	 Capital-augmenting Makes existing equipment more productive without changing the worker's tasks. AMBIGUOUS	 Automation Substitutes software for tasks that workers previously performed; commodifies expertise. NOT PRO-WORKER
---	--	--	--	---

Source: Acemoglu, Autor, and Johnson (2026).

THREE DISTINCTIONS THE TERM DEPENDS ON

Task displacement is not job loss

An occupation can lose tasks and retain or gain employment if demand for its remaining tasks grows.

Wage levels are not labor's share

Wages can rise for some workers while the share of income accruing to labor falls.

Worker outcomes are not total welfare

A technology can be good for patients and payers while being bad for a specific workforce.

The workers whose outcomes count are those in the occupations that perform the affected tasks. The horizon is the next decade, over which current products are deployed at scale.

02 · THREE ILLUSTRATIVE DEPLOYMENTS

One occupation, three futures

To make the framework concrete, I adapt the device AAJ use for delivery drivers to community health workers (CHWs), describing three hypothetical AI deployments that a health system might adopt to extend CHW reach. The scenarios are illustrative. They vary both the underlying capabilities and the deployment design, so they do not isolate the effect of workflow design alone. Their purpose is to show that the range of plausible labor consequences is wide even for a single occupation with a single stated goal.

AUTOMATION

Community Care Automator

This tool scripts each patient interaction through a branching decision tree, generates daily task lists from geolocation, electronic health record (EHR) integration, and algorithmic risk scores, and monitors worker compliance through GPS tracking and time-stamped task completion. The design premise is that outreach can be performed by any worker who can follow the script, and the developer's stated long-term aim is to replace human outreach with patient-facing chatbots and automated messaging. In AAJ terms the tool automates the judgment component of CHW work and commodifies the relational expertise that distinguishes the occupation.

LABOR-AUGMENTING

Community Care Assistant

This tool uses a large language model to provide real-time decision support when a CHW encounters an unfamiliar clinical or social situation. A CHW visiting a patient newly diagnosed with heart failure can query the tool about medication side effects, local food resources that stock low-sodium options, or eligibility for a transportation benefit. Predictive analytics flag patients whose missed appointments, unfilled prescriptions, or emergency department visits suggest emerging risk. The tool surfaces information as input to the CHW's own judgment and does not script the interaction. In AAJ terms it is labor-augmenting, and its effect on the value of CHW expertise is ambiguous: it raises productivity but may narrow the gap between experienced and novice workers.

NEW TASK-CREATING

Community Care Enhancement

This tool equips CHWs with FDA-cleared point-of-care diagnostic devices, such as a dermatoscope attachment with computer-vision triage of skin lesions, a portable electrocardiogram monitor that detects atrial fibrillation, or a retinal camera with an autonomous diabetic retinopathy algorithm, each generating a preliminary reading that a specialist confirms. The CHW does not replace the specialist. The CHW performs screening in community settings where, as a matter of hypothesis, no one was previously performing it because the equipment and interpretive capacity existed only in clinical facilities. In AAJ terms this is new task-creating for the CHW occupation and expertise-leveling with respect to the specialist, and whether it is pro-worker in net depends on whether it reaches unmet need or redistributes existing demand.

The three scenarios share natural language processing, predictive analytics, and computer vision. What differs is the role assigned to human expertise in the workflow and the choice of which tasks to give the worker. The pro-worker character of health AI is therefore, in principle, a design and policy variable rather than a technological given. The empirical question is whether the market is producing the full range of these trajectories or concentrating investment in a subset. The environmental scan in Part III is a first, descriptive answer to that question for commercially deployed products.



03

The environmental scan

How 132 technologies were found, screened, and classified, what the distribution looks like, and how much the headline result moves under alternative coding.

03 • DESIGN

A product-level inventory under a pre-specified protocol

I conducted a structured environmental scan under a protocol written before the search began (Protocol version 3.0, March 2026; summarized in the Supplementary Methods). The unit of analysis is a named AI-enabled product or platform. Each row in the dataset is one product from one vendor, identified by company, product name, domain, subdomain, source URL, and a factual summary of what the product does, its deployment scale, and its regulatory status. Where a vendor markets several AI products, I coded the flagship health AI product identified by the vendor's own materials, so that each vendor appears once. This choice makes a large diversified company and a single-product startup count equally, which is appropriate for describing the composition of the product landscape but not for describing its market weight. The scan is not adoption-weighted, and readers should interpret proportions as shares of products, not shares of workers, revenue, or encounters.

Technologies were eligible if AI or machine learning is a core functional component rather than a marketing label; the product is deployed or commercially available in the US health system; the product has reached one of three maturity states; and the product operates in at least one of the five study domains. The three maturity states are distinct and are recorded separately in the dataset: commercial deployment with documented US customers, FDA marketing authorization, or, for drug-discovery platforms only, at least one asset in Phase II or later with a defined commercialization pathway. The third state was included because AI drug-discovery companies rarely sell software to health systems and would otherwise be excluded from a scan of the life sciences domain. General-purpose AI platforms, internal government tools that are not commercially deployed, internal pharmaceutical company platforms, consumer wellness applications without clinical evidence or FDA involvement, EHR systems without a distinct AI module, and pre-commercial research prototypes were excluded. The consequence of these criteria for the public health domain is discussed in Section 4.2.3.

Search, screening, and classification

1

FDA list of AI-enabled medical devices, with specialty searches

2

Market intelligence: KLAS, Rock Health, CB Insights, investor disclosures

3

Trade press and targeted searches across 12 administrative subdomains

4

PubMed for empirical studies of tools with US clinical deployments

5

J.P. Morgan 2026 and HIMSS 2025 announcements; venture rounds above \$50M

6

Three LLM research platforms to improve retrieval coverage (not validation)

The search proceeded in six phases between March 22 and March 28, 2026: the FDA list of AI-enabled medical devices; market intelligence sources (KLAS Research, Rock Health, CB Insights, investor disclosures); trade publications and targeted searches for each of 12 administrative subdomains; PubMed for empirical studies of AI tools with US clinical deployments; proceedings and announcements from the 2026 J.P. Morgan Healthcare Conference and the 2025 HIMSS conference, together with venture rounds above \$50 million; and, in the sixth phase, repeated searches using three large language model research platforms prompted with the same inclusion criteria and domain taxonomy. The research agents were run under the protocol's supplementary instructions for LLM-based research agents (Section 9), which require that every factual claim be attributable to a retrieved source and that the LLM-generated dataset be reviewed by a human domain expert. They contributed candidate products and draft dataset entries, including provisional classifications; I do not describe their output as validation of the classifications. Every candidate surfaced by any phase was screened by me against the eligibility criteria using its primary source, and the source URL recorded for each included product is a company product page or an FDA database entry. Candidates were added iteratively until successive searches within a subdomain yielded no new eligible products; this stopping rule is operational, not a demonstration of completeness, and the scan should be read as a large purposive sample of the landscape rather than a census.

I reviewed every draft entry and provisional classification produced by the research agents under the protocol's human-review step: I verified source URLs, checked classifications against the decision tree, reviewed the labor and wage assessments for plausibility, and revised entries where needed. The final classifications are mine. The protocol also specified a second analyst who would independently classify at least 20% of technologies so that inter-rater agreement could be reported. That step was not completed, and no agreement statistic is available. This is the most important methodological limitation of the scan, and the classifications should be read as one trained health economist's structured application of the AAJ decision rules, assisted by LLM research agents and documented row by row in the dataset so that others can audit or re-code them.

For each product I recorded the anticipated direction of effect on three dimensions specified by AAJ, each coded as positive, negative, ambiguous, or approximately zero, with a written rationale: the effect on labor productivity in the affected tasks, the effect on the value of human expertise in the affected occupations, and the effect on labor's share of income from the affected activity. The AAJ category was then assigned by a decision tree that uses those three codes together with a counterfactual question: what did workers do before the product, what does the product substitute for, and what additional human work does it make possible. Because the three directional codes are inputs to the category assignment, the within-category uniformity of those codes reported below is a property of the coding rule, not independent evidence about the categories. I report it to make the rule transparent, not as a finding.

The distinction between new work and more work (Acemoglu and Restrepo 2019) governed the new task-creating category. A product was coded as new task-creating only when the dataset rationale identified a human task that is new to the occupation or setting, not merely a new technical workflow for an existing task. I did not require that the task be new to the economy, and I did not assess net employment effects, because product characteristics cannot establish either. Where a product plausibly fit more than one category, I assigned the predominant anticipated consequence and recorded the alternative in the rationale. To bound the sensitivity of the headline results to the two most contestable coding decisions, I report in Section 3.4 what the distribution would look like under plausible alternative assignments. These are arithmetic recalculations on the same dataset, not re-coding.

03 • RESULTS


Automation leads the count, but the domains tell the story

Of 132 technologies, 55 (41.7%) were classified as automating existing tasks, 39 (29.5%) as new task-creating, 22 (16.7%) as expertise-leveling, 14 (10.6%) as labor-augmenting, and 2 (1.5%) as capital-augmenting. Table 1 defines each category and gives counts, principal domains, illustrative products, and the occupations most often named in the dataset rationales. Figure 1 shows the within-domain distribution. The full technology-level dataset, with a unique identifier for each product, is provided as Supplementary Table S1 and Supplementary Data S1.

EXHIBIT 3

Four in ten technologies were classified as automation; three in ten as new task-creating.

Distribution of 132 AI-enabled US health technologies by AAJ pro-worker category




Source: Infectious Economics environmental scan, March 2026 (Appendix B).

The domain pattern is the central descriptive result. Provider administration (26 products) and payer administration (14 products) together contain 40 technologies, of which 39 (97.5%) were classified as automation, 1 (2.5%) as labor-augmenting, and none as new task-creating or expertise-leveling. Clinical care (56 products) is the only domain with entries in all five categories: 15 expertise-leveling (26.8%), 14 new task-creating (25.0%), 13 automation (23.2%), 12 labor-augmenting (21.4%), and 2 capital-augmenting (3.6%). Life sciences (21 products) is almost entirely new task-creating (20, or 95.2%). Public health (15 products) contains 7 expertise-leveling (46.7%), 5 new task-creating (33.3%), and 3 automation (20.0%). Pooling the three health-producing domains, 61 of 92 technologies (66.3%) were classified as new task-creating or expertise-leveling, compared with 0 of 40 in the two administrative domains. All denominators are stated in the text because the paper alternates between shares of all 132 technologies, shares of a domain, and shares of a category.

EXHIBIT 4

Seen as shares of each domain, the same counts show two administrative bars that are almost entirely automation.

Within-domain share of technologies by category; labels show count and percentage for segments of 6% or more



Source: Infectious Economics environmental scan (Appendix B). Unlabeled segments: capital-augmenting in clinical care (2 of 56); labor-augmenting in provider administration (1 of 26) and life sciences (1 of 21).

03 • CATEGORY REFERENCE

The five categories of pro-worker AI, as applied to the scan

Counts are shares of all 132 technologies. Illustrative products are examples, not exhaustive lists. Occupations are those most frequently named by Standard Occupational Classification code in the dataset rationales. Status follows Acemoglu, Autor, and Johnson (2026) and refers to anticipated effects, not observed outcomes.

● Labor-augmenting 14 (10.6%)

Makes workers more productive at tasks they already perform; tasks stay with the worker.

PRINCIPAL DOMAINS

Clinical care (12); provider administration (1); life sciences (1)

ILLUSTRATIVE PRODUCTS

AI second readers and triage in imaging and pathology (Aidoc, Lunit, Paige, PathAI, RapidAI); capacity management (LeanTaaS)

OCCUPATIONS MOST OFTEN NAMED

Radiologists; pathologists; cardiologists

AAJ STATUS

Ambiguous

● Capital-augmenting 2 (1.5%)

Makes existing equipment more productive without changing the worker's tasks.

PRINCIPAL DOMAINS

Clinical care (2)

ILLUSTRATIVE PRODUCTS

Deep learning scan acceleration (Subtle Medical); device parameter optimization (ResMed)

OCCUPATIONS MOST OFTEN NAMED

Radiologic technologists

AAJ STATUS

Ambiguous

● Automation 55 (41.7%)

Substitutes software for tasks that workers previously performed.

PRINCIPAL DOMAINS

Provider administration (25); payer administration (14); clinical care (13); public health workforce (3)

ILLUSTRATIVE PRODUCTS

Revenue cycle and coding (R1 RCM, Waystar, CodaMetrix); prior authorization (Cohere Health); claims (Availity, Cotiviti); ambient documentation (Abridge, Ambience, Microsoft/Nuance, Suki); staffing marketplaces (Aya, ShiftMed)

OCCUPATIONS MOST OFTEN NAMED

Medical records specialists; medical scribes and transcriptionists; billing and posting clerks; customer service representatives; claims examiners

AAJ STATUS

Not pro-worker

● New task-creating

39 (29.5%)

Generates demand for human expertise in tasks that did not previously exist in the occupation or setting.

PRINCIPAL DOMAINS

Life sciences (20); clinical care (14); public health (5)

ILLUSTRATIVE PRODUCTS

Generative drug design (Insilico, Recursion, Generate); genomics workflows (Illumina, SOPHiA); real-world evidence (Flatiron, Truveta); coronary CT physiology (HeartFlow); stroke coordination (Viz.ai); surgical robotics (Intuitive); surveillance (BlueDot)

OCCUPATIONS MOST OFTEN NAMED

Data scientists; medical and biological scientists; computational chemists; genetic counselors; stroke coordinators; epidemiologists

AAJ STATUS

Unambiguously pro-worker

● Expertise-leveling

22 (16.7%)

Enables less-expert workers, or workers in new settings, to perform tasks that previously required greater expertise.

PRINCIPAL DOMAINS

Clinical care (15); public health (7)

ILLUSTRATIVE PRODUCTS

Autonomous retinal screening (Digital Diagnostics, Eyenuk); point-of-care dermatology (DermaSensor); guided ultrasound (Butterfly, Ultrasight); AI-guided primary care and virtual care (K Health, Doctronic, Teladoc); digital therapeutics (Sword, Hinge)

OCCUPATIONS MOST OFTEN NAMED

Nurse practitioners; physician assistants; primary care physicians; community health workers


AAJ STATUS

Ambiguous

EXHIBIT 5

Medical records specialists are the occupation most often named, and almost always in an automation rationale.

Twelve occupations most frequently named in the labor-impact rationales of the 132 technologies, colored by the AAJ category of the technology naming them



Source: Infectious Economics environmental scan; occupations identified by Standard Occupational Classification code in the dataset rationales. A technology counts once per occupation it names.

Sensitivity of the headline distribution to alternative coding

Two coding decisions carry most of the weight in the results, and both are contestable. The first is the classification of ambient clinical documentation tools as automation. Nine products in the clinical documentation subdomain (Abridge, Ambience Healthcare, Augmedix, Doximity, Freed AI, Microsoft/Nuance, NextGen Healthcare, RXNT, Suki AI) were coded as automation because they substitute for the work of medical scribes and transcriptionists. A reviewer who weighted the effect on physicians would code them as labor-augmenting. Under that alternative, automation falls from 55 to 46 technologies (34.8%) and labor-augmenting rises from 14 to 23 (17.4%). The second is the classification of 20 life sciences products as new task-creating. Drug discovery, trial design, and real-world evidence generation existed before these products, and a skeptic would code them as labor-augmenting for scientists who already perform those activities. Under that alternative, new task-creating falls from 39 to 19 technologies (14.4%), and the pooled health-producing share of new task-creating or

expertise-leveling falls from 61 of 92 (66.3%) to 41 of 92 (44.6%). Under both alternatives together, automation is 46 of 132 (34.8%) and new task-creating is 19 of 132 (14.4%). The administrative result is unaffected by either alternative: 39 of 40 administrative technologies remain classified as automation. The domain divergence is therefore robust to the two most plausible re-codings, while the absolute size of the new task-creating category is not.

EXHIBIT 6

The administrative finding holds under alternative coding. Other shares depend on it.

Share of technologies under the base classification and under the alternative coding: nine ambient documentation products and twenty life sciences products moved to labor-augmenting

	BASE CLASSIFICATION		ALTERNATIVE CODING
All products classified as automation	41.7%	→	34.8%
All products classified as new task-creating	29.5%	→	14.4%
Health-producing domains: new task-creating or expertise-leveling	66.3%	→	44.6%
Administrative domains: automation	97.5%	=	97.5%

Source: Arithmetic recalculation on the scan dataset (Appendix B); no technology was re-coded.

The domain divergence is robust to the two most plausible re-codings. The absolute size of the new task-creating category is not.

03 · CATEGORY-LEVEL DESCRIPTION

What sits inside each category

Labor-augmenting 14 of 132, 10.6%

Twelve of the 14 are in clinical care, one is in provider administration (a capacity-management platform), and one is in life sciences (a data-integration platform). Most of the clinical entries are AI-assisted second readers and triage tools in diagnostic imaging and pathology that increase clinician throughput on interpretive tasks the clinician still performs. Under the coding rule, all 14 raise labor productivity; for 12 the effect on the value of expertise is coded ambiguous because clinician judgment remains central while the marginal value of routine pattern recognition declines. Whether such tools benefit patients depends on the payment context, discussed in Part IV: under fee-for-service, additional imaging throughput can serve unmet need or generate low-value volume, and product characteristics cannot distinguish the two.

Capital-augmenting 2 of 132, 1.5%

Two products make existing equipment more productive: deep learning reconstruction that shortens MRI and PET acquisition times on installed scanners, and machine learning that optimizes positive airway pressure device settings. The near-absence of this category is itself descriptive of health AI: almost every product in the scan acts on the labor input rather than the capital input, consistent with a sector in which the binding constraints are clinician time and administrative processing rather than equipment capacity.

Automation 55 of 132, 41.7%

Thirty-nine of the 55 (70.9%) are in provider or payer administration: revenue cycle management and coding, prior authorization, claims adjudication, credentialing, patient billing, and collections. Thirteen are in clinical care, of which nine are ambient documentation tools, two are imaging-report or cardiac-monitoring automation, one is an operating-room sensing platform, and one is a patient-facing conversational agent. Three are in public health, all of them staffing and workforce marketplaces coded within the health workforce subdomain. The occupations most often named in the rationales are medical records specialists, medical transcriptionists and scribes, billing and posting clerks, and customer service representatives. The automation category therefore spans clinical and workforce tools as well as administrative ones, although administrative tools dominate it. Ambient scribes illustrate the dual character of clinical automation: they raise physician productivity and may reduce documentation burden and burnout, while displacing a scribe workforce that itself emerged only in the past decade in response to EHR documentation demands. The net effect on workers as a group cannot be inferred from the classification and requires outcome data.

New task-creating 39 of 132, 29.5%

Twenty of the 39 are in life sciences: generative design of proteins and small molecules, clinical trial optimization, precision genomics workflows, and real-world evidence platforms. Fourteen are in clinical care, including non-invasive coronary physiology assessment from CT angiography, AI-directed stroke care coordination, remote patient monitoring, and surgical robotics. Five are in public health: infectious disease surveillance and social-needs analytics. The claim I make for this category is deliberately narrow. Each product was coded as creating a task that is new to the occupation or setting in which it is deployed, as documented in the rationale. A computational chemist running foundation-model simulations performs work that requires a combination of machine learning, structural biology, and pharmacology expertise that the occupation did not previously require. That is a new task in the AAJ sense. It does not establish that the technology creates net employment, and early evidence on whether AI-discovered molecules succeed in clinical trials is only beginning to accumulate (Jayatunga et al. 2024). Biopharmaceutical research and development is a large enough activity, with global investment estimated at \$276 billion in 2021 (Chandra et al. 2024), that even modest changes in its task composition matter for the scientific workforce.

Expertise-leveling 22 of 132, 16.7%

Fifteen are in clinical care and seven in public health: FDA-cleared point-of-care diagnostics that let primary care clinicians screen for conditions that previously required specialist referral, AI-guided ultrasound for non-sonographers, decision support that extends specialist judgment to generalists, digital therapeutics, and virtual care platforms, including one platform that delivers autonomous AI consultations under physician oversight. The occupations most often named are nurse practitioners, physician assistants, primary care physicians, and community health workers. For all 22, the effect on the value of expertise and on labor's share is coded ambiguous, reflecting the offsetting effects the AAJ framework predicts: entrants gain capability while incumbents face erosion of scarcity rents. Roughly 92 million Americans lived in a designated primary care Health Professional Shortage Area at the end of 2025 (KFF 2026), so the access stakes are large. Whether the trajectory is pro-worker depends on whether expanded capacity meets previously unmet need, as the Grossman (1972) model of health demand would predict when access barriers fall, or redistributes existing demand from specialists to lower-cost substitutes.



04

Three domains, three economic logics

Five institutional features decide who captures AI's gains. They play out differently in the friction-cost layer, in clinical care, and in public health.

04 • INSTITUTIONAL FEATURES

Five features that shape who captures the gains

The scan documents what the market is building. To interpret why the pattern differs by domain, and who is likely to capture the resulting gains, I draw on five institutional features of US health care markets. I describe them as institutional features rather than as established market failures, because some of them generate inefficiency, some respond to it, and several do both. For each, I state the condition under which the proposed mechanism operates.

1 Information asymmetry Sustains demand for clinical judgment; patient-facing AI must first solve trust and accountability.	2 Consolidation and employer power Where few systems employ most local workers, gains are retained as margin rather than paid as wages.	3 Insurance tied to employment Administrative savings relieve the wage squeeze only if they pass through to premiums.	4 Payment design Fee-for-service can serve unmet need or generate low-value volume; capitation can redirect time or retain savings.	5 Licensure and scope of practice Constrain AI-augmented practice even where the technology is capable; incumbents hold concentrated interests.
--	--	--	--	--

First, information asymmetry between clinicians and patients (Arrow 1963) sustains demand for clinical judgment as long as patients cannot verify the quality of the advice they receive. This does not by itself establish that clinical judgment will resist automation. What it establishes is that a patient-facing AI product must solve a trust and accountability problem, through regulation, liability, or professional oversight, before it can substitute for a clinician rather than augment one. The scan's clinical results, where augmenting and leveling products outnumber automating ones, are consistent with that constraint but do not prove it. Second, provider consolidation has raised prices in output markets (Gaynor and Town 2011; Brot-Goldberg et al. 2024), and the same consolidation can confer employer market power in local labor markets for health workers. These are different markets and require distinct evidence. The relevant condition for this paper is labor market concentration: where a small number of systems employ most of the local clinical or administrative workforce, productivity gains are more likely to be retained as margin than passed through as wages, because workers' outside options are constrained. Third, the linkage of health insurance to employment (Gruber 2000) means that administrative savings, if passed through to premiums, would relieve the compensation squeeze that rising health costs impose on workers across the economy. The condition is pass-through, which depends on competition among insurers and employers. Fourth, provider payment design mediates how a throughput gain is used (Chernew 2020). Under fee-for-service, additional capacity can serve unmet need or generate low-value volume, and which one occurs depends on demand responsiveness and utilization safeguards. Under capitation, freed clinician time can be redirected to higher-value care or the savings can be retained without improving care or working conditions. Payment reform therefore changes incentives but does not by itself guarantee beneficial deployment. Fifth, licensure and scope-of-practice regulation constrain AI-augmented practice even where the technology is capable (Dranove 2011;

Nicholson and Propper 2011), creating a political economy in which incumbent professions hold concentrated interests in exclusivity while the beneficiaries of expanded access are diffuse.

04 • THE FRICTION-COST DOMAIN

The friction-cost paradox

\$43B

spent by US hospitals in 2025
collecting payments from insurers for
care already delivered (AHA 2026)

34.2%

of US health expenditures went to
administration in 2017, against
17.0% in Canada (Himmelstein et al.
2020)

\$265.6B

estimated annual waste from
administrative complexity, the largest
category in Shrank et al. (2019)

The tasks automated in provider and payer administration are rule-bound functions defined by billing codes, contract terms, and coverage policies. They are algorithmically reducible in the way the AAJ framework predicts. Much of the workforce performing them is an artifact of institutional complexity: the American Hospital Association estimates that hospitals spent \$43 billion in 2025 attempting to collect payments from insurers for care already delivered, and that the average hospital employs about 64 administrative and billing staff dedicated to these functions, roughly 6.5% of hospital employment (AHA 2026). Administrative costs accounted for an estimated 34.2% of US health expenditures in 2017 compared with 17.0% in Canada (Himmelstein et al. 2020), and US administrative spending as a share of total health spending is roughly three to eight times that of other high-income countries (Papanicolas et al. 2018). Within Shrank et al. (2019), administrative complexity is the single largest category of estimated waste at \$265.6 billion per year. The friction-cost layer targeted by administrative AI is therefore a product of the multi-payer architecture rather than a feature of care delivery.

This creates what I call the friction-cost paradox: the workers most exposed to automation perform work that exists because of institutional complexity, and reducing that work would improve efficiency at the system level while concentrating losses on a specific workforce. The paradox requires two refinements. First, administrative work is not uniformly waste. It is useful to distinguish three kinds: work that is unnecessary and could be eliminated by standardization or by removing the requirement that generates it, such as duplicative documentation requests; work that is necessary but performed inefficiently, such as credentialing, fraud prevention, scheduling, and appeals, which protect patients and payers even when their current implementation is costly; and work that principally redistributes payments between payers and providers, such as denial and appeal cycles, where one side's cost is the other side's revenue. Automation is the appropriate response mainly to the second kind. For the first kind, eliminating the requirement dominates automating compliance with it. For the third kind, AI on both sides of a dispute can intensify the contest, raising the volume of coded claims, denials, and appeals without reducing the resources the system devotes to them. A private financial return to a revenue-cycle product is not the same as a reduction in social resource use.

04 • DISTRIBUTIONAL CONSEQUENCES

The transition falls on identifiable workforces

92.3%

Medical records specialists

84.5%

Billing and posting clerks

89.6%

Receptionists and information clerks

Women as a share of employed workers in each occupation, 2025 annual averages. Source: BLS (2026), Current Population Survey, Table 11. Percentages describe the occupations, not observed AI displacement.

Second, the distributional case for managing the transition is strong, and the evidence for it should be occupation specific. The occupations most often named in the automation rationales are predominantly female: in 2025, women were 92.3% of employed medical records specialists, 84.5% of billing and posting clerks, and 89.6% of receptionists and information clerks (BLS 2026). The scan does not document the educational attainment or geographic concentration of these workers, and those characteristics should be established from occupational data before they are used to design transition programs. Where these workers are employed by consolidated systems with few local alternatives, the fifth feature above implies that gains will be retained rather than shared. Welfare gains from administrative automation are therefore conditional: on preserving the functions that protect patients and payers, on accounting for implementation cost, error rates, and burden shifted to patients and clinicians, and on the losses borne by displaced workers.

Task displacement is not job loss. Wage levels are not labor's share. Worker outcomes are not total welfare.

EXHIBIT 7

Administrative work is not uniformly waste; automation is the right response to only one of three kinds.

Unnecessary work

Duplicative documentation requests and requirements that could be eliminated by standardization.

Best response: remove the requirement

Necessary work done inefficiently

Credentialing, fraud prevention, scheduling, appeals. These protect patients and payers even when costly.

Best response: automate carefully, preserve the function

Redistributive work

Denial and appeal cycles where one side's cost is the other's revenue. AI on both sides can intensify the contest.

Best response: standardize rules, not the arms race

Source: Infectious Economics analysis.

A private financial return to a revenue-cycle product is not the same as a reduction in social resource use.

04 • THE CLINICAL CARE DOMAIN

Clinical AI splits along two trajectories

Clinical care is the largest domain in the scan and the only one with entries in all five categories. The heterogeneity is consistent with an economic structure in which diagnostic interpretation, treatment planning, and patient communication require discretion and contextual reasoning, which the AAJ framework identifies as the setting where human-AI collaboration is most productive. The scan shows clinical AI splitting along two trajectories: one that augments or extends expertise (second readers, point-of-care diagnostics, decision support), and one that automates documentation. Which trajectory absorbs more deployment, and whether either serves patients, depends on the conditions described in Section 4.1: demand responsiveness and utilization safeguards under fee-for-service, and the use to which freed time is put under capitation. Payment reform is a lever that changes the odds, not a prerequisite that guarantees the outcome.

Payment reform is a lever that changes the odds, not a prerequisite that guarantees the outcome.

04 • THE PUBLIC HEALTH DOMAIN

Few commercial tools serve core public health functions

\$14.2B

raised by US digital health startups in 2025; AI-enabled companies captured 54% (Rock Health 2026)

80,000

additional full-time staff needed by state and local health departments (Leider et al. 2023)

20 to 47%

less pay in government public health for management, computer, and scientific occupations than in the private sector (Krasna et al. 2024)

Public health contributes 15 of 132 technologies, and only 5 were coded as new task-creating public health tools in the narrow sense of surveillance or social-needs analytics. The scan's design limits what can be concluded from this. The eligibility criteria emphasize commercial deployment, FDA authorization, and venture funding, and they exclude internal government tools and academic systems. Those choices are likely to undercount public health innovation, much of which is built inside health departments, universities, and nonprofit consortia and procured rather than sold. The public health domain in the scan also includes population health analytics and workforce staffing platforms, which complicates interpretation. The correct statement of the result is therefore that among commercially oriented AI products meeting the scan's criteria, few address core public health functions.

Separate evidence is consistent with the hypothesis that private capital underinvests in those functions. Core public health outputs such as outbreak detection and population-level prevention are non-rivalrous and largely non-excludable, with no fee-for-service reimbursement to support per-unit pricing. Venture funding responds to that structure: US digital health startups raised \$14.2 billion in 2025, AI-enabled companies captured 54% of it, and the most-funded value propositions were clinical and non-clinical workflow tools, which together drew 39% of all funding (Rock Health 2026). The institutions that would deploy public health AI face a workforce constraint as well: state and local health departments were estimated to need roughly 80,000 additional full-time staff (Leider et al. 2023), and management, computer, and scientific occupations in government public health pay 20% to 47% less than the same occupations in the private sector (Krasna et al. 2024). Federal workforce reductions and grant terminations in 2025 have further reduced capacity (McKillop et al. 2025). Public health is funded jointly by federal, state, and local governments, with the federal government the principal funder of many core functions through CDC and related agencies. The public-good character of these services varies: outbreak detection is close to a pure public good, whereas community screening yields private benefits that can support reimbursement. If public purchasers do not create demand for tools that serve the former, the private market is unlikely to build them.



05

Policy can change the trajectory

If the domain pattern reflects market structure and payment incentives, policy can change the trajectory. Each lever names a responsible actor, a mechanism, implementation requirements, and an evaluation criterion. Support should follow demonstrated benefit, not novelty.

-
- 01 Align coding, coverage, and payment with beneficial deployment
 - 02 Make administrative AI savings visible
 - 03 Test new scopes of practice under supervision
 - 04 Buy public health tools, and fund the capacity to use them
 - 05 Give displaced workers a path and a voice

05 • POLICY LEVERS

Support should follow demonstrated benefit, not novelty

If the domain pattern reflects market structure and payment incentives, policy can change the trajectory. The five levers below name a responsible actor, a mechanism, implementation requirements, and evaluation criteria. Each faces identifiable opposition: providers who benefit from volume incentives, consolidated systems that capture administrative savings as margin, specialty associations that defend scope exclusivity, and fiscal pressure on public health budgets. Naming the opposition is a reason to sequence the levers, not to abandon them. A principle runs through all five: support should follow demonstrated patient benefit, access, job quality, and efficient resource use, not novelty. A new task is not valuable because it is new.

Naming the opposition is a reason to sequence the levers, not to abandon them. A new task is not valuable because it is new.

05 · LEVER 1 OF 5

Align coding, coverage, and payment with beneficial deployment

CMS · Medicare Administrative Contractors · state Medicaid agencies

Three decisions are often conflated: whether a service has a billing code, whether a payer covers it, and what it pays. Non-invasive coronary physiology assessment illustrates the sequence. A Category I CPT code exists (75580) and Medicare contractors cover the service under a local coverage determination with documentation requirements (CMS 2024); the remaining gaps are coverage variation across payers and payment levels that do not reward the downstream avoided catheterizations. For expertise-leveling diagnostics, CMS should test payment models in which screening performed in federally qualified health centers and rural clinics is reimbursed at levels that reflect access gains, conditional on documented follow-up. For labor-augmenting tools in imaging, throughput gains should be paired with appropriate-use criteria and outcome measurement so that added capacity serves unmet need rather than low-value volume.

Coding

Does a service have a billing code?

Coverage

Is the service covered for this population?

Payment

What payment supports valuable deployment?

EVALUATION CRITERION

Change in access, diagnostic yield, and downstream utilization in covered versus uncovered markets.

05 • LEVER 2 OF 5

Require reporting of administrative AI savings and their allocation

CMS for Medicare Advantage • ACA plans • state insurance regulators • hospital cost-report authorities

Payers and health systems that deploy administrative AI at scale should report the estimated savings, the method used to estimate them, and the allocation among premium reductions, clinical staffing, and retained margin. Savings can be estimated from cost-report line items for administrative labor and vendor contracts, audited against the same reports, and compared over time within the organization. Disclosure alone will not change distribution. Its purpose is to make the allocation visible to regulators who set medical loss ratio requirements and to employers who purchase coverage, and to generate the data needed to test the friction-cost paradox empirically.

Estimate

Savings from cost-report labor lines and vendor contracts

Audit

Checked against the same reports over time

Disclose

Allocation to premiums, staffing, or margin made visible

EVALUATION CRITERION

Share of reported savings passed through to premiums or clinical staffing, by market concentration.

05 • LEVER 3 OF 5

Modernize workforce regulation through supervised pilots

state legislatures • licensing boards • with HRSA support

States should authorize time-limited pilots in which nurse practitioners, physician assistants, and community health workers use FDA-cleared AI diagnostics under defined protocols, with a named supervising clinician accountable for the program, mandatory adverse-event reporting, and prospective collection of safety and outcome data. Pilots should specify in advance the outcomes that would justify permanent scope changes. Federal support should expand training pathways in clinical informatics and applied data science for public health.

Protocol

Defined tasks and FDA-cleared tools

Supervision

A named accountable clinician and adverse-event reporting

Outcomes

Pre-specified criteria for permanent scope change

EVALUATION CRITERION

Diagnostic accuracy, adverse events, and access relative to usual referral pathways.

05 · LEVER 4 OF 5

Create public demand for public health AI

CDC · ARPA-H · state health departments

The scan cannot establish investment levels, but the public-good character of core public health functions is a sufficient reason for public purchasers to specify and procure the tools the private market has little incentive to build. Two things must be separated: acquiring software, and funding the staff, maintenance, data infrastructure, and implementation that make it useful. Procurement without sustained operating support will produce shelfware. The HITECH Act's use of financial incentives to accelerate EHR adoption shows that public purchasing can redirect a technology trajectory; it also shows the cost of neglecting usability and interoperability.

Software

Procure tools that serve public health needs

Implementation

Data infrastructure and integration

Operating capacity

Retain staff and sustain maintenance and use

EVALUATION CRITERION

Time to outbreak detection, coverage of surveillance systems, and retention of the analytical workforce in participating agencies.

05 • LEVER 5 OF 5

Fund intra-sector transition support and give workers a voice in deployment

Department of Labor • CMS • health system employers • unions

Administrative workers displaced by automation hold knowledge of health care operations, coding rules, and payer requirements that is relevant to adjacent roles in clinical informatics, quality measurement, and AI oversight. Transferability is not automatic: many of those roles require additional credentials, retraining takes time and money, openings are fewer than the workers displaced, and wages differ. Transition programs modeled on Trade Adjustment Assistance but oriented to intra-sector reallocation should fund credentialing, provide income support during retraining, and be evaluated on re-employment and earnings. Given the role of bargaining power in determining who captures gains, workers and their representatives should participate in procurement and deployment decisions, alongside the transparency and retraining measures above.

Credentials

Fund the training adjacent roles require

Income

Support during retraining, evaluated on earnings

Voice

Workers at the table for procurement and deployment

EVALUATION CRITERION

Re-employment rate, earnings replacement, and time to re-employment for displaced workers.

06 • LIMITATIONS

What this scan can and cannot support

The limitations of this study follow from its design. The scan is a purposive inventory of commercially oriented products, not a census, and it excludes government, academic, and nonprofit systems; its coverage of public health is therefore incomplete by construction. The unit of analysis is the product, unweighted by adoption, so proportions describe the composition of the landscape rather than its labor market weight. Draft entries and provisional classifications were produced by large language model research agents under a pre-specified protocol, and all final classifications were reviewed and made by a single analyst; the planned inter-rater reliability check by a second analyst was not completed. The three directional assessments that feed the classification are anticipated consequences, not observed outcomes, and their within-category consistency is a property of the coding rule. The sensitivity analysis in Section 3.4 bounds the two most contestable coding choices but does not exhaust them. Finally, the economic interpretation in Part IV is a set of hypotheses about mechanisms. Testing them requires linking product deployment to occupational employment and earnings data over time, which is the natural next study.

The classifications and the interpretation in this report are my own judgments, made from publicly available information: company product pages, FDA database entries, press coverage, and published studies, as of March 2026. I did not test any product, contact any vendor, or use non-public information. A category describes a product's anticipated effect on work under the framework, not its quality, safety, or clinical effectiveness, and a vendor may reasonably disagree with a classification.

06 • CONCLUSION

A design variable, not a technological given

Within a scan of 132 commercially deployed AI-enabled health technologies, the anticipated labor consequences differ sharply by domain. Administrative products were almost all classified as automating existing tasks; clinical, life sciences, and public health products were mostly classified as creating new tasks or extending expertise to new workers and settings. The divergence is robust to the two most plausible alternative codings, while the absolute size of the new task-creating category is not.

I have argued that this pattern is consistent with the underlying economics of each domain, and that who captures the gains from health AI will depend on payment incentives, labor market concentration, insurance pass-through, and workforce regulation rather than on the technology alone. The friction-cost paradox is the central tension: the workers most exposed to automation perform work that exists because of institutional complexity, and the welfare case for reducing that work is conditional on preserving the functions that protect patients, on avoiding a coding and denial arms race, and on managing the losses of a predominantly female clerical workforce. The same capabilities that could script and monitor a community health worker could augment her judgment or equip her to screen for disease where no one now does. Which of these the US health system gets will be decided by how CMS structures coverage and payment, whether administrative savings are made visible, whether scope-of-practice rules adapt under supervision, whether public purchasers create demand for tools that serve public goods, and whether displaced workers are given a path and a voice. The framework and dataset presented here are offered as a starting point for making those choices with evidence.

The same capabilities that could script and monitor a community health worker could augment her judgment or equip her to screen for disease where no one now does. Which of these the US health system gets will be decided by policy choices made in the next several years.

ABOUT THE AUTHOR



Blythe Adamson, PhD, MPH

Blythe Adamson helps health technology companies, health systems, and public agencies fill gaps in health and economic evidence to show which products and policies work and what they are worth. She is a health economist and epidemiologist and the founder of Infectious Economics in New York. Her work spans cost-effectiveness analysis, real-world evidence strategy, and the economics of artificial intelligence in health care. She served in the White House in 2020 and has led research at Flatiron Health and Doctrionic.

More about her work, publications, and this report at blytheadamson.com.

Disclosures

Funding. This work received no external funding and was supported by Infectious Economics, LLC.

Conflicts of interest. The author was employed by Flatiron Health before this study and received fees from Doctrionic Inc. after the scan and classifications were completed in March 2026. Both companies appear in the scan: Flatiron's OncoEMR and real-world data platform is coded as new task-creating, and Doctrionic's AI doctor platform is coded as expertise-leveling. Neither company reviewed or influenced the classifications. The author is also an affiliate professor at the University of Washington in the Comparative Health Outcomes, Policy, and Economics (CHOICE) Institute. The classifications and interpretation are the author's own judgments from publicly available information, and the views expressed are the author's alone and do not represent those of any employer, client, or institution with which the author is affiliated, or of the authors of the framework.

Use of artificial intelligence. The author used large language model tools for editing, including prose editing, structural reorganization, and formatting. All analytical judgments, classifications, policy recommendations, and substantive content are the author's own. Large language model research agents were also used in the environmental scan workflow, under the protocol's supplementary instructions for LLM-based research agents, to search for candidate products and to draft dataset entries and provisional classifications. The author reviewed every LLM-generated entry, verified sources, checked classifications against the decision tree, and revised entries where needed. The author takes full responsibility for the accuracy and integrity of the work.

Acknowledgments. The author thanks Simon Johnson for conversations in early 2026 that shaped the design of this project. Special thanks to Gabriela Cejas, predoctoral researcher at MIT, for reviewing a draft of this paper and providing feedback. Any errors are the author's own.

Data availability. The full technology-level dataset (Supplementary Data S1) and the protocol are available with this paper.

REFERENCES

Sources

1. Acemoglu, D., Autor, D., & Johnson, S. (2026). Building pro-worker artificial intelligence. NBER Working Paper No. 34854. Cambridge, MA: National Bureau of Economic Research. Written for the Hamilton Project, Brookings Institution. <https://doi.org/10.3386/w34854>
2. Acemoglu, D., & Restrepo, P. (2018). Artificial intelligence, automation, and work. NBER Working Paper No. 24196. Cambridge, MA: National Bureau of Economic Research. <https://doi.org/10.3386/w24196>
3. Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3-30. <https://doi.org/10.1257/jep.33.2.3>
4. American Hospital Association (AHA). (2026). Costs of caring: Challenges facing America's hospitals as they care for patients in 2026. Chicago, IL: American Hospital Association, March 2026. <https://www.aha.org/costsofcaring>
5. Arrow, K. J. (1963). Uncertainty and the welfare economics of medical care. *American Economic Review*, 53(5), 941-973.
6. Autor, D. (2024). Applying AI to rebuild middle class jobs. NBER Working Paper No. 32140. Cambridge, MA: National Bureau of Economic Research. <https://doi.org/10.3386/w32140>
7. Brot-Goldberg, Z., Cooper, Z., Craig, S. V., Klarinet, L., Lurie, I., & Miller, C. (2024). Who pays for rising health care prices? Evidence from hospital mergers. NBER Working Paper No. 32613. Cambridge, MA: National Bureau of Economic Research. <https://doi.org/10.3386/w32613>
8. Centers for Medicare and Medicaid Services (CMS). (2024). Billing and coding: Non-invasive fractional flow reserve (FFR) for ischemic heart disease. Medicare Coverage Database, Article A58473. <https://www.cms.gov/medicare-coverage-database/view/article.aspx?articleId=58473>
9. Centers for Medicare and Medicaid Services (CMS). (2025). National health expenditures 2024 highlights. Baltimore, MD: CMS Office of the Actuary. <https://www.cms.gov/data-research/statistics-trends-and-reports/national-health-expenditure-data/historical>
10. Chandra, A., Drum, J., Daly, M., Mirsberger, H., Spare, S., Neumann, U., Martin, S., & Kirson, N. (2024). Comprehensive measurement of biopharmaceutical R&D investment. *Nature Reviews Drug Discovery*, 23(9), 652-653. <https://doi.org/10.1038/d41573-024-00131-2>
11. Chernenov, M. E. (2020). The role of market forces in U.S. health care. *New England Journal of Medicine*, 383(15), 1401-1404. <https://doi.org/10.1056/NEJMp2006097>
12. Dranove, D. (2011). Health care markets, regulators, and certifiers. In M. V. Pauly, T. G. McGuire, & P. P. Barros (Eds.), *Handbook of Health Economics* (Vol. 2, pp. 639-690). Amsterdam: Elsevier. <https://doi.org/10.1016/B978-0-444-53592-4.00010-4>
13. Gaynor, M., & Town, R. J. (2011). Competition in health care markets. In M. V. Pauly, T. G. McGuire, & P. P. Barros (Eds.), *Handbook of Health Economics* (Vol. 2, pp. 499-637). Amsterdam: Elsevier. <https://doi.org/10.1016/B978-0-444-53592-4.00009-8>
14. Grossman, M. (1972). On the concept of health capital and the demand for health. *Journal of Political Economy*, 80(2), 223-255. <https://doi.org/10.1086/259880>
15. Gruber, J. (2000). Health insurance and the labor market. In A. J. Culyer & J. P. Newhouse (Eds.), *Handbook of Health Economics* (Vol. 1, pp. 645-706). Amsterdam: Elsevier. [https://doi.org/10.1016/S1574-0064\(00\)80171-7](https://doi.org/10.1016/S1574-0064(00)80171-7)
16. Himmelstein, D. U., Campbell, T., & Woolhandler, S. (2020). Health care administrative costs in the United States and Canada, 2017. *Annals of Internal Medicine*, 172(2), 134-142. <https://doi.org/10.7326/M19-2818>
17. Jayatunga, M. K. P., Ayers, M., Bruens, L., Jayanth, D., & Meier, C. (2024). How successful are AI-discovered drugs in clinical trials? A first analysis and emerging lessons. *Drug Discovery Today*, 29(6), 104009. <https://doi.org/10.1016/j.drudis.2024.104009>
18. Keehan, S. P., Madison, A. J., Poisal, J. A., Cuckler, G. A., Smith, S. D., Sisko, A. M., Fiore, J. A., & Rennie, K. E. (2025). National health expenditure projections, 2024-33: Despite insurance coverage declines, health to grow as share of GDP. *Health Affairs*, 44(7), 776-787. <https://doi.org/10.1377/hlthaff.2025.00545>
19. KFF. (2026). Primary care health professional shortage areas (HPSAs), as of December 31, 2025. San Francisco, CA: KFF, based on HRSA Bureau of Health Workforce data. <https://www.kff.org/other/state-indicator/primary-care-health-professional-shortage-areas-hpsas/>
20. Krasna, H., Venkataraman, M., & Patino, I. (2024). Salary disparities in public health occupations: Analysis of federal data, 2021-2022. *American Journal of Public Health*, 114(3), 329-339. <https://doi.org/10.2105/AJPH.2023.307512>
21. Leider, J. P., McCullough, J. M., Singh, S. R., Sieger, A., Robins, M., Fisher, J. S., & Castrucci, B. C. (2023). Staffing up and sustaining the public health workforce. *Journal of Public Health Management and Practice*, 29(3), E100-E107. <https://doi.org/10.1097/PHH.0000000000001614>

22. McKillop, M., Leiberman, D., & Farberman, R. K. (2025). Public health infrastructure in crisis: HHS workforce cuts, reorganizations, and funding reductions: Impacts and solutions. Washington, DC: Trust for America's Health. https://www.tfah.org/wp-content/uploads/2025/09/TFAH_Funding_Report2025_FINAL.pdf
23. Nicholson, S., & Propper, C. (2011). Medical workforce. In M. V. Pauly, T. G. McGuire, & P. P. Barros (Eds.), *Handbook of Health Economics* (Vol. 2, pp. 873-925). Amsterdam: Elsevier. <https://doi.org/10.1016/B978-0-444-53592-4.00014-1>
24. Papanicolas, I., Woskie, L. R., & Jha, A. K. (2018). Health care spending in the United States and other high-income countries. *JAMA*, 319(10), 1024-1039. <https://doi.org/10.1001/jama.2018.1150>
25. Rock Health. (2026). 2025 year-end digital health funding overview: A tale of two markets. San Francisco, CA: Rock Health, January 2026. <https://rockhealth.com/insights/2025-year-end-digital-health-funding-overview-a-tale-of-two-markets/>
26. Shrank, W. H., Rogstad, T. L., & Parekh, N. (2019). Waste in the US health care system: Estimated costs and potential for savings. *JAMA*, 322(15), 1501-1509. <https://doi.org/10.1001/jama.2019.13978>
27. U.S. Bureau of Labor Statistics (BLS). (2023). Healthcare occupations in 2022. Spotlight on Statistics, December 2023. <https://www.bls.gov/spotlight/2023/healthcare-occupations-in-2022/>
28. U.S. Bureau of Labor Statistics (BLS). (2026). Labor force statistics from the Current Population Survey, Table 11: Employed persons by detailed occupation, sex, race, and Hispanic or Latino ethnicity, 2025 annual averages. <https://www.bls.gov/cps/cpsaat11.htm>

APPENDIX A

Supplementary methods

Protocol. The scan followed a pre-specified protocol (Environmental Scan and Pro-Worker Classification of AI-Enabled Health Technologies in the United States, Protocol version 3.0, March 2026), which defines the domain taxonomy, eligibility criteria, data fields, the three directional assessments, the classification decision tree, and quality assurance procedures. The protocol's inter-rater reliability step (independent classification of at least 20% of technologies by a second analyst, with a target Cohen's kappa of at least 0.60) was not completed; all screening and classification was performed by the author.

Domain taxonomy. Domain A, clinical care delivery (seven subdomains: diagnostic imaging and radiology; pathology and laboratory medicine; clinical documentation and EHR workflow; clinical decision support and predictive analytics; surgical and procedural care; remote monitoring, telehealth, and virtual care; digital therapeutics and behavioral health). Domain B, provider-side administration, friction-cost layer (five subdomains: revenue cycle management and coding; patient access, scheduling, and eligibility; patient billing and collections; provider credentialing and enrollment; hospital operations and capacity management). Domain C, payer-side administration, friction-cost layer (six subdomains: claims processing and adjudication; prior authorization and utilization management; payment integrity and fraud detection; risk adjustment and quality measurement; member services and benefits administration; pharmacy benefit management). Domain D, life sciences and biomedical R&D (four subdomains: drug discovery and computational chemistry; clinical trial design and recruitment; genomics and precision medicine; real-world evidence and health data platforms). Domain E, public health, population health, and health workforce (three subdomains: disease screening and surveillance; population health analytics and social determinants; health workforce and staffing).

Search phases and sources. Phase 1: the FDA list of AI-enabled medical devices, with targeted searches by specialty (radiology, cardiology, ophthalmology, dermatology, neurology). Phase 2: KLAS Research reports, Rock Health funding databases, CB Insights healthcare AI rankings, and company investor disclosures. Phase 3: healthcare IT trade publications, revenue cycle management industry reports, and targeted searches for companies operating in each of 12 administrative subdomains (claims processing, prior authorization, utilization management, payment integrity, fraud detection, risk adjustment, medical coding, patient billing, provider credentialing, pharmacy benefit management, member services, hospital operations). Phase 4: PubMed searches for empirical studies documenting AI tools with US clinical deployments. Phase 5: announcements from the 2026 J.P. Morgan Healthcare Conference and the 2025 HIMSS Global Health Conference, and venture rounds above \$50 million. Phase 6: searches using three large language model research platforms (Claude Opus 4.6 with extended thinking, Anthropic; ChatGPT 5.4 Deep Research, OpenAI; Gemini 3 Thinking Pro, Google DeepMind), each prompted with the same eligibility criteria and domain taxonomy on March 26-28, 2026. Each platform surfaced candidates the others did not. The platforms were run as research agents under the protocol's Section 9 instructions and also drafted dataset entries and provisional classifications. All candidates from all phases were screened by the author against the eligibility criteria using the primary source recorded in the dataset, and all LLM-generated entries were reviewed by the author before inclusion. The search was conducted March 22-28, 2026.

Screening counts. The dataset records the 132 products that met all eligibility criteria. Counts of candidates screened and excluded at each phase, with reasons, were not retained in a form that permits a PRISMA-style flow diagram; this is a limitation of the scan. Deduplication was performed at the vendor level, with one flagship health AI product per vendor. Version 8 of the dataset (March 30, 2026) reflects two additions made after an internal consistency review of version 7.

Eligibility. Included: (a) AI or machine learning is a core functional component; (b) the product is deployed or commercially available in the US health system; (c) the product is in one of three maturity states, recorded in the summary field: commercial deployment with documented US customers, FDA marketing authorization, or (drug discovery only) at least one asset in Phase II or later with a defined commercialization pathway; (d) the product operates in at least one study domain. Excluded: general-purpose AI platforms; internal government tools not commercially deployed; internal pharmaceutical company platforms not independently marketed; consumer wellness applications without clinical evidence, FDA involvement, or clinical system integration; EHR systems without a distinct AI module; research or pre-commercial prototypes.

Classification. For each product the author coded the anticipated direction of effect on labor productivity, the value of human expertise, and labor's share of income (positive, negative, ambiguous, or approximately zero), each with a written rationale, and recorded the affected occupations by SOC code, the evidence tier supporting the deployment and performance claims (Tier 1, prospective clinical study or trial; Tier 2, retrospective or observational study; Tier 3, vendor claim or case study), and a narrative assessment of anticipated effects on labor demand and wages. The AAJ category was assigned by the protocol decision tree. Genuine ambiguity was recorded in the rationale and the predominant anticipated consequence was assigned.

APPENDIX B

Technology-level dataset (n = 132)

Columns: ID; company; product; domain (A clinical care, B provider administration, C payer administration, D life sciences, E public health); subdomain; anticipated effect on labor productivity (LP), value of human expertise (VE), and labor's share of income (LS), coded + positive, – negative, +/- ambiguous, ≈0 approximately zero; AAJ category (1 labor-augmenting, 2 capital-augmenting, 3 automation, 4 new task-creating, 5 expertise-leveling). The full record for each product, including summary, source URL, rationales, affected occupations, and evidence tier, is available for download, under a CC BY 4.0 license, at blytheadamson.com/pro-worker-ai-in-health.

ID	Company	Product	Dom.	Sub.	LP	VE	LS	AAJ
T001	Aidoc	aiOS Platform	A	A1	+	+/-	≈0	1
T002	AliveCor	Kardia 12L / KAI 12L	A	A1	+	+/-	+/-	5
T003	Cleerly	Cleerly Coronary Artery Disease AI	A	A1	+	+	+	4
T004	EchoGo	EchoGo Heart Failure 2.0	A	A1	+	+/-	≈0	1
T005	GE HealthCare	Edison AI / Caption Health	A	A1	+	+/-	+/-	5
T006	HeartFlow	HeartFlow FFRCT	A	A1	+	+	+	4
T007	Lunit	Lunit INSIGHT (Mammography / Chest)	A	A1	+	+/-	+/-	1
T008	Nanox (Zebra Medical Vision)	Zebra AI Imaging Suite	A	A1	+	+/-	+	4
T009	Overjet	Overjet CBCT Assist	A	A1	+	+/-	+/-	5
T010	Rad AI	Rad AI Reporting / Impressions / Continuity	A	A1	+	+/-	–	3
T011	RapidAI	RapidAI Stroke/Aneurysm Platform	A	A1	+	+/-	≈0	1
T012	Siemens Healthineers	AI-Rad Companion / syngo.via	A	A1	+	+/-	≈0	1
T013	Subtle Medical	SubtleMR / SubtlePET / SubtleGAD	A	A1	+	+/-	≈0	2
T014	Viz.ai	Viz Care Coordination Platform	A	A1	+	+	+	4
T015	iCAD	ProFound AI Software V3.0	A	A1	+	+/-	≈0	1
T016	iRhythm Technologies	Zio Monitor / Deep-Learned AI	A	A1	+	+/-	–	3
T017	Ibex Medical Analytics	Galen Second Read	A	A2	+	+	≈0	1
T018	Paige	Paige Prostate / PanCancer Detect	A	A2	+	+	≈0	1
T019	PathAI	AI Sight Platform	A	A2	+	+/-	≈0	1
T020	Abridge	Abridge Ambient AI	A	A3	+	+/-	–	3
T021	Ambience Healthcare	Ambience AutoScribe	A	A3	+	+/-	–	3
T022	Augmedix (Commure)	Augmedix Live / Augmedix Go	A	A3	+	+/-	–	3
T023	Doximity	Doximity Scribe / DoxGPT / PeerCheck	A	A3	+	–	–	3
T024	Freed AI	Freed AI Scribe	A	A3	+	–	–	3
T025	Microsoft/Nuance	DAX Copilot (Dragon Ambient eXperience)	A	A3	+	+/-	–	3
T026	NextGen Healthcare	NextGen Ambient Assist	A	A3	+	+/-	–	3
T027	Pomelo Care	Pomelo Care Virtual Maternity & Women's Health Platform	A	A3	+	+	+	4
T028	RXNT	Ambient IQ	A	A3	+	–	–	3

ID	Company	Product	Dom.	Sub.	LP	VE	LS	AAJ
T029	Suki AI	Suki Ambient Intelligence	A	A3	+	+/-	-	3
T030	CURRENT Health (Best Buy Health)	AI Remote Patient Monitoring	A	A4	+	+	+	4
T031	Ceribell	Delirium Monitor System	A	A4	+	+/-	+/-	5
T032	Eko Health	Eko AI Cardiac Screening	A	A4	+	+/-	+/-	5
T033	Epic Systems	Epic AI / Cognitive Computing	A	A4	+	+/-	≈0	1
T034	Health Catalyst	Health Catalyst Data Platform	A	A4	+	+	+	4
T035	Innovaccer	Innovaccer Health Cloud	A	A4	+	+	+	4
T036	Natus Medical / Persyst	BrainWatch / Persyst 15	A	A4	+	+/-	+/-	5
T037	OpenEvidence	OpenEvidence AI Clinical Decision Platform	A	A4	+	+	+	4
T038	ResMed	Personalized Therapy Comfort Settings (PTCS)	A	A4	+	+/-	≈0	2
T039	Ultrasight	UltraSight Guidance	A	A4	+	+/-	+/-	5
T040	Artisight	Artisight Smart Hospital Platform	A	A5	+	-	-	3
T041	Caresyntax	AI Surgical Analytics Platform	A	A5	+	+	+	4
T042	Intuitive Surgical	da Vinci 5 / Ion / Case Insights AI	A	A5	+	+	+	4
T043	Iterative Health	SKOUT System	A	A5	+	+/-	≈0	1
T044	Johnson & Johnson	Ottava / MONARCH Platform	A	A5	+	+	+	4
T045	Medtronic	Hugo RAS / GI Genius / MiniMed AI	A	A5	+	+/-	+/-	1
T046	Zimmer Biomet	ROSA AI Knee Replacement	A	A5	+	+	+	4
T047	Doctronic	Doctronic AI Doctor Platform	A	A6	+	+/-	+/-	5
T048	Hippocratic AI	Polaris LLM / AI Health Agents	A	A6	+	-	-	3
T049	K Health	K Health AI Virtual Primary Care	A	A6	+	+/-	+/-	5
T050	Teladoc Health	AI-Enhanced Virtual Care	A	A6	+	+/-	+/-	5
T051	Click Therapeutics (formerly Pear)	Click Suite / reSET / Somryst	A	A7	+	+/-	+/-	5
T052	Grow Therapy	AI Care Companion	A	A7	+	+	+	4
T053	Hinge Health	Hinge Health Digital MSK Platform	A	A7	+	+/-	+/-	5
T054	Sword Health	Sword AI Phoenix Platform	A	A7	+	+/-	+/-	5
T055	Virta Health	Virta AI-Powered Diabetes Reversal	A	A7	+	+/-	+/-	5
T056	Woebot Health	Woebot	A	A7	+	+/-	+/-	5
T057	AKASA	AKASA GenAI RCM Platform	B	B1	+	-	-	3
T058	Aspirion	DocIQ / ContractIQ	B	B1	+	-	-	3
T059	CodaMetrix	CodaMetrix Autonomous Coding	B	B1	+	-	-	3
T060	CombineHealth	CombineHealth AI Revenue Cycle Platform	B	B1	+	-	-	3
T061	Commure	Commure Pro (Intake-to-Payment AI)	B	B1	+	-	-	3
T062	ENTER Health	ENTER RCM / CTRL ENTER	B	B1	+	-	-	3
T063	Fathom	Fathom Autonomous Medical Coding	B	B1	+	-	-	3
T064	Iodine Software (Waystar)	Iodine AwareCDI	B	B1	+	-	-	3
T065	Nym Health	Nym Autonomous Coding Engine	B	B1	+	-	-	3

ID	Company	Product	Dom.	Sub.	LP	VE	LS	AAJ
T066	R1 RCM	R1 AI-Powered RCM Platform	B	B1	+	-	-	3
T067	RapidClaims.AI	RapidCode / RapidScrub / RapidCDI	B	B1	+	-	-	3
T068	SmarterDx	SmarterPrebill / SmarterDenials	B	B1	+	-	-	3
T069	Clearstep	Clearstep Smart Care Routing	B	B2	+	+/-	-	3
T070	Experian Health	Patient Access Curator / AI Advantage	B	B2	+	-	-	3
T071	Hyro	Hyro AI Agents Platform	B	B2	+	-	-	3
T072	Notable Health	Notable AI Platform	B	B2	+	-	-	3
T073	Tennr	Tennr AI Referral & Document Processing	B	B2	+	-	-	3
T074	Cedar	Cedar Intelligence / Kora Voice Agent	B	B3	+	-	-	3
T075	Collectly	Collectly AI Patient Billing / Voice Agent	B	B3	+	-	-	3
T076	PayZen	PayZen AI Patient Financing Platform	B	B3	+	-	-	3
T077	Assured	Assured AI Provider Network Management	B	B4	+	-	-	3
T078	Medallion	Medallion AI Credentialing Platform	B	B4	+	-	-	3
T079	Verifiable	Verifiable AI Credentialing	B	B4	+	-	-	3
T080	sympplr	sympplr Provider Credentialing AI	B	B4	+	-	-	3
T081	LeanTaaS	iQueue Platform	B	B5	+	+/-	≈0	1
T082	Qventus	Qventus AI Operations	B	B5	+	+/-	-	3
T083	Availity	AuthAI / Intelligent UM	C	C1	+	-	-	3
T084	Cotiviti	Cotiviti AI Claims Analytics	C	C1	+	-	-	3
T085	HealthEdge	HealthEdge AI-Powered Payer Platform	C	C1	+	-	-	3
T086	Optum (UnitedHealth)	Optum Real / AI Claims Platform	C	C1	+	-	-	3
T087	Banjo Health	BanjoPA / BanjoAppeals	C	C2	+	-	-	3
T088	Cohere Health	Cohere UM Suite / Align / Validate	C	C2	+	-	-	3
T089	Shift Technology	Shift Claims Fraud Detection AI	C	C3	+	-	-	3
T090	Zelis	Zelis Payment Integrity / Network Management AI	C	C3	+	-	-	3
T091	Apixio (Clover Health)	Apixio AI Risk Adjustment / HCC Coding	C	C4	+	-	-	3
T092	Episource (Cotiviti)	Episource AI Risk Adjustment	C	C4	+	-	-	3
T093	HDAI	HDAI AI Claims Summaries	C	C5	+	-	-	3
T094	Waystar (absorbed Olive AI)	Waystar AI-Powered Platform	C	C5	+	-	-	3
T095	CoverMyMeds (McKesson)	CoverMyMeds ePA Platform	C	C6	+	-	-	3
T096	Surescripts	Surescripts AI Prior Auth for Medications	C	C6	+	-	-	3
T097	Absci Corporation	Integrated Drug & Target Discovery	D	D1	+	+	+	4
T098	Generate Biomedicines	Generative Biology Platform	D	D1	+	+	+	4
T099	Inductive Bio	Inductive Bio Compass Platform / Beacon ADMET Models	D	D1	+	+	+	4
T100	Insilico Medicine	Pharma.AI / ISM001-055	D	D1	+	+	+	4
T101	Insitro	Insitro Platform / COAST-SS	D	D1	+	+	+	4
T102	NVIDIA	BioNeMo / Clara / MONAI	D	D1	+	+	+	4

ID	Company	Product	Dom.	Sub.	LP	VE	LS	AAJ
T103	Recursion Pharmaceuticals	Recursion OS / Exscientia Platform	D	D1	+	+	+	4
T104	Schrödinger	Schrödinger Platform	D	D1	+	+	+	4
T105	Paradigm Health	AI Clinical Trial Recruitment	D	D2	+	+	+	4
T106	Tempus AI	Tempus Lens / TIME Trial Matching / David AI	D	D2	+	+	+	4
T107	Unlearn.AI	Unlearn TwinRCT	D	D2	+	+	+	4
T108	Foundation Medicine (Roche)	FoundationOne CDx / FoundationOne Liquid CDx	D	D3	+	+	+	4
T109	Illumina	DRAGEN AI / Connected Analytics	D	D3	+	+	+	4
T110	SOPHiA GENETICS	SOPHiA DDM Platform / Digital Twins	D	D3	+	+	+	4
T111	Clarify Health	Clarify AI Analytics	D	D4	+	+	+	4
T112	Flatiron Health (Roche)	OncoEMR / RWD Platform	D	D4	+	+	+	4
T113	IQVIA	IQVIA Analytics Research Accelerator	D	D4	+	+	+	4
T114	Komodo Health	Komodo MapEnhance	D	D4	+	+	+	4
T115	Mayo Clinic	Mayo Clinic Platform (MCP)	D	D4	+	+	+	4
T116	Palantir Technologies	Palantir Foundry for Healthcare	D	D4	+	+/-	+/-	1
T117	Truveta	Truveta Studio / Language Model for Medicine	D	D4	+	+	+	4
T118	AEYE Health	AEYE Diagnostic Screening (AEYE-DS)	E	E1	+	+/-	+/-	5
T119	BlueDot	BlueDot Intelligence Solution	E	E1	+	+	+	4
T120	Butterfly Network	Butterfly iQ+ / AI Guidance	E	E1	+	+/-	+/-	5
T121	DermaSensor	DermaSensor Skin Cancer Detection	E	E1	+	+/-	+/-	5
T122	Digital Diagnostics	LumineticsCore (formerly IDx-DR)	E	E1	+	+/-	+/-	5
T123	Eyenuk	EyeArt AI Eye Screening	E	E1	+	+/-	+/-	5
T124	Qure.ai	qXR / qER AI Diagnostics	E	E1	+	+/-	+/-	5
T125	Yale School of Public Health	PopHIVE (Population Health Information and Visualization Exchange)	E	E1	+	+	+	4
T126	CareMessage	CareMessage Wellness Copilot / AI Response Assistant	E	E2	+	+/-	+/-	5
T127	Cityblock Health	Cityblock AI Care Platform	E	E2	+	+	+	4
T128	Socially Determined	SocialScape Platform	E	E2	+	+	+	4
T129	Unite Us	Unite Us AI Social Care Coordination	E	E2	+	+	+	4
T130	Aya Healthcare	Aya Workforce AI	E	E3	+	+/-	-	3
T131	Incredible Health	AI Nurse Matching Platform	E	E3	+	+/-	-	3
T132	ShiftMed	ShiftMed Workforce AI Suite	E	E3	+	-	-	3

The future of health AI is a policy choice.

How productivity gains are shared matters as much as how they are produced.



INFECTIOUS ECONOMICS

Economic evidence for health technology

New York, NY · blytheadamson.com · blythe@infectious-economics.com