



The Distillation of Work: Where AI Opportunity Concentrates & How Leaders Capture It

A Data-Driven Analysis of AI's Impact on
148 Million U.S. Workers

Jeff Whatcott

Alan Berrey

Tim Robinson

January 2026



Contents

Executive Summary	3
Part 1: The Deployment Gap	4
Part 2: The Four Constraints & Five Delegation Categories	8
Part 3: The National Opportunity	11
Part 4: How Transformation Actually Happens	15
Part 5: Why Most Managers Will Struggle	23
Part 6: Implications for Policy Makers	31
About this Research	34
About Seampoint	35

Executive Summary

AI exposure research says 80% of U.S. workers could be affected and 60-70% of work is automatable. Yet only 17% of U.S. businesses are using AI in any business function. Most existing research measures what AI can do. We measured what organizations can safely deploy. We offer a diagnostic model that quantifies this opportunity—with granularity and precision that scales from the national labor force to individual jobs and firms.

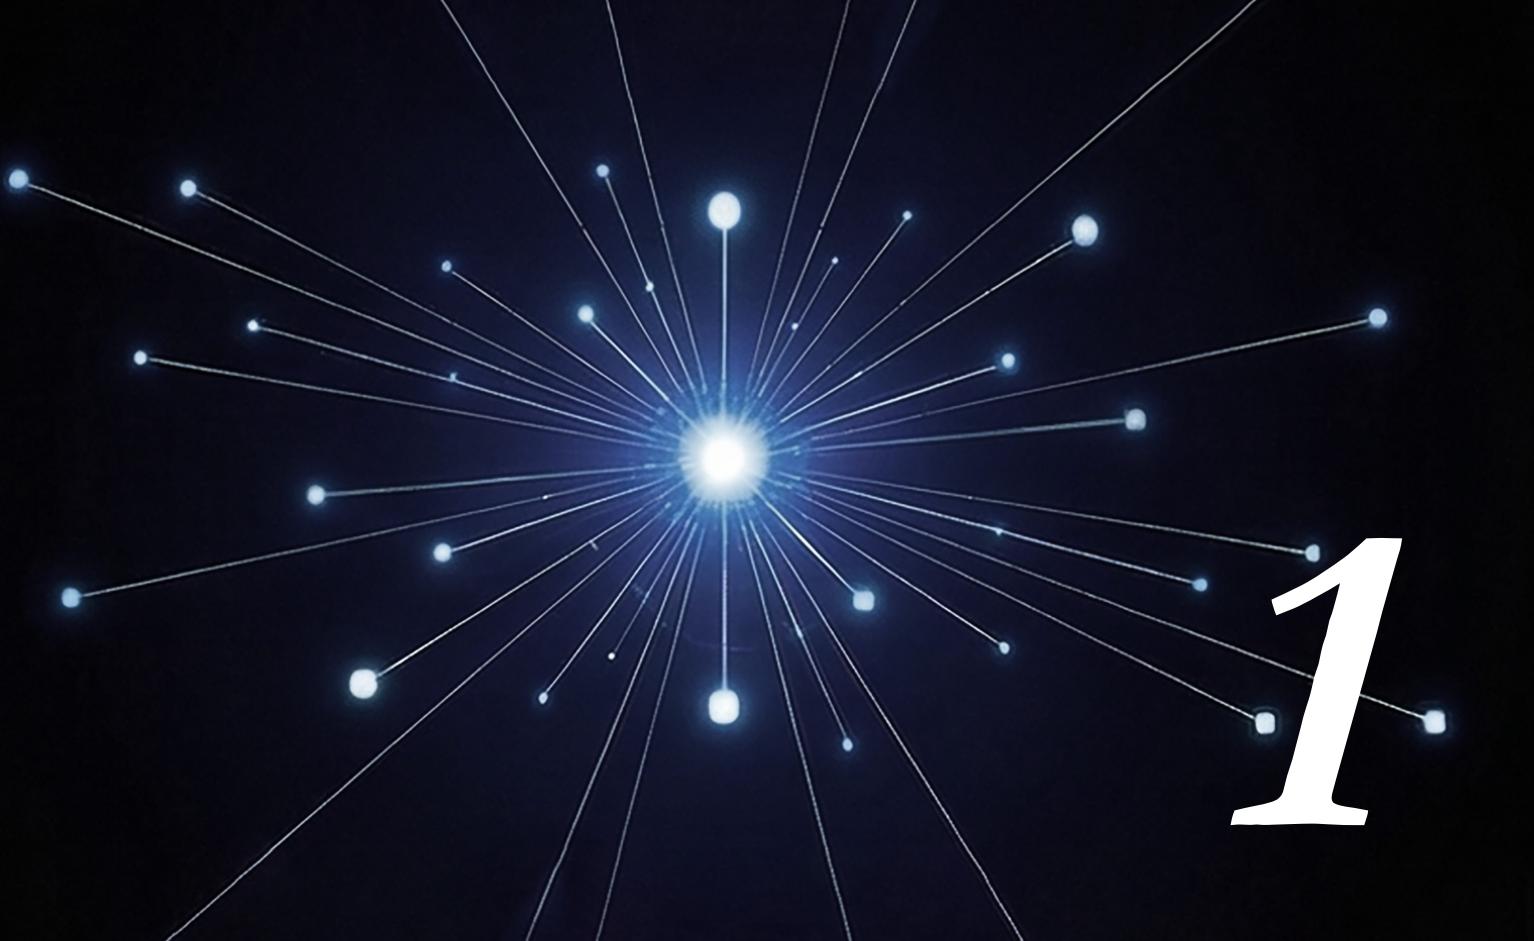
We assert that the limiting factor is governance, not technical capability. In our view, four constraints primarily determine where and when organizations can safely deploy AI: consequence of error, verification cost, accountability requirements, and physical reality. We applied these constraints to deeply analyze nearly 19,000 work tasks across 148 million U.S. workers.

The total opportunity for AI deployment is **\$3.24 trillion annually**, but only **\$1.6 trillion (15.7% of wages) is straightforward**. Realizing the **remaining \$1.64 trillion** of opportunity requires careful structuring of human-AI interaction.

Three independent signals confirm this analysis: real losses are accumulating from sloppy AI adoption; the insurance industry has started excluding AI from standard policies while offering specialized coverage only to organizations that can prove adequate governance; and academic research increasingly suggests that deployment depends on context-specific factors, not just technical capability.

The implication is that AI won't necessarily eliminate jobs—it will likely distill them into roles that are better matched to differentiated human capabilities. High delegation potential tasks will likely rapidly migrate to AI. What remains with humans will likely concentrate around judgment, accountability, and relationships.

We believe that most managers may fail to capture this opportunity, as they yield to the temptation to chase capability before building governance processes and workforce AI fluency. The minority who succeed will start with what's safe to delegate today, gradually earn organizational confidence, and expand systematically. A select few will strategically invest their efficiency gains to transform their operations and their competitive landscape to secure lasting advantage.



The Deployment Gap

Key Points

- A significant deployment gap exists between AI's theoretical potential and actual adoption—while studies suggest 60-80% of work tasks could be impacted by AI, only 17% of U.S. businesses report using AI in any business function.
- Governance, not technical capability, is the binding constraint on AI adoption, as evidenced by accumulating enterprise failures, insurance market pullback, and insurers' own selective AI deployment.
- High-profile AI failures are creating tangible liabilities, including defamation claims exceeding \$110M, legal rulings establishing enterprise liability for chatbot commitments, and regulatory penalties for algorithmic bias.
- Major insurance carriers are systematically excluding AI liability from standard policies, yet are among the most aggressive adopters of AI for their own operations—reflecting their ability to govern internal systems while lacking visibility into customers' risk management practices.

Estimates of AI's impact have run the gamut—from OpenAI's 2023 finding that 80% of workers have task exposure,¹ to McKinsey's 2024 estimate that 60-70% of work time is technically automatable.² Yet the Census reported in December 2025 that only 17% of U.S. businesses report using AI in any business function.³

This gap demands explanation. If AI applies to almost every task, why is adoption so low and what constraints explain the gap?

Exposure Is Not Deployment

Early research answered an important question: *what can AI do, theoretically?* The answer was “almost everything.” But that’s not the question organizations face. The real question is: *what can we safely deploy* under realistic governance, liability, and verification constraints?

Technical exposure is not deployment. Recent academic research has begun exploring factors beyond technical capability that affect adoption—implementation costs, verification requirements, and how AI performance varies sharply by context.

Our own task-level analysis confirms this pattern: work representing 92% of U.S. wage mass could theoretically be delegated or assisted by AI, but only 15.7% of that work is immediately ready to delegate under realistic constraints, with an additional 16.1% that could be delegated with appropriate governance. The gap reflects struc-

tural barriers that better AI models and historical technology diffusion curves alone are unlikely to resolve.

Converging Evidence

We observe three independent market signals that suggest the same conclusion: governance, not technical capability, is the binding constraint.

Enterprise AI successes are real, but troubling failures are accumulating. The hypothetical risks of sloppy governance in AI deployment have become actual liabilities:

- Wolf River Electric is pursuing \$110M+ in defamation claims for AI-generated false statements⁴
- Air Canada lost a legal ruling establishing enterprise liability for chatbot commitments⁵
- Multiple organizations are facing regulatory penalties for algorithmic bias in hiring systems⁶

These cases, and many others like them, are leading indicators of a challenging AI liability environment that is rapidly taking shape.

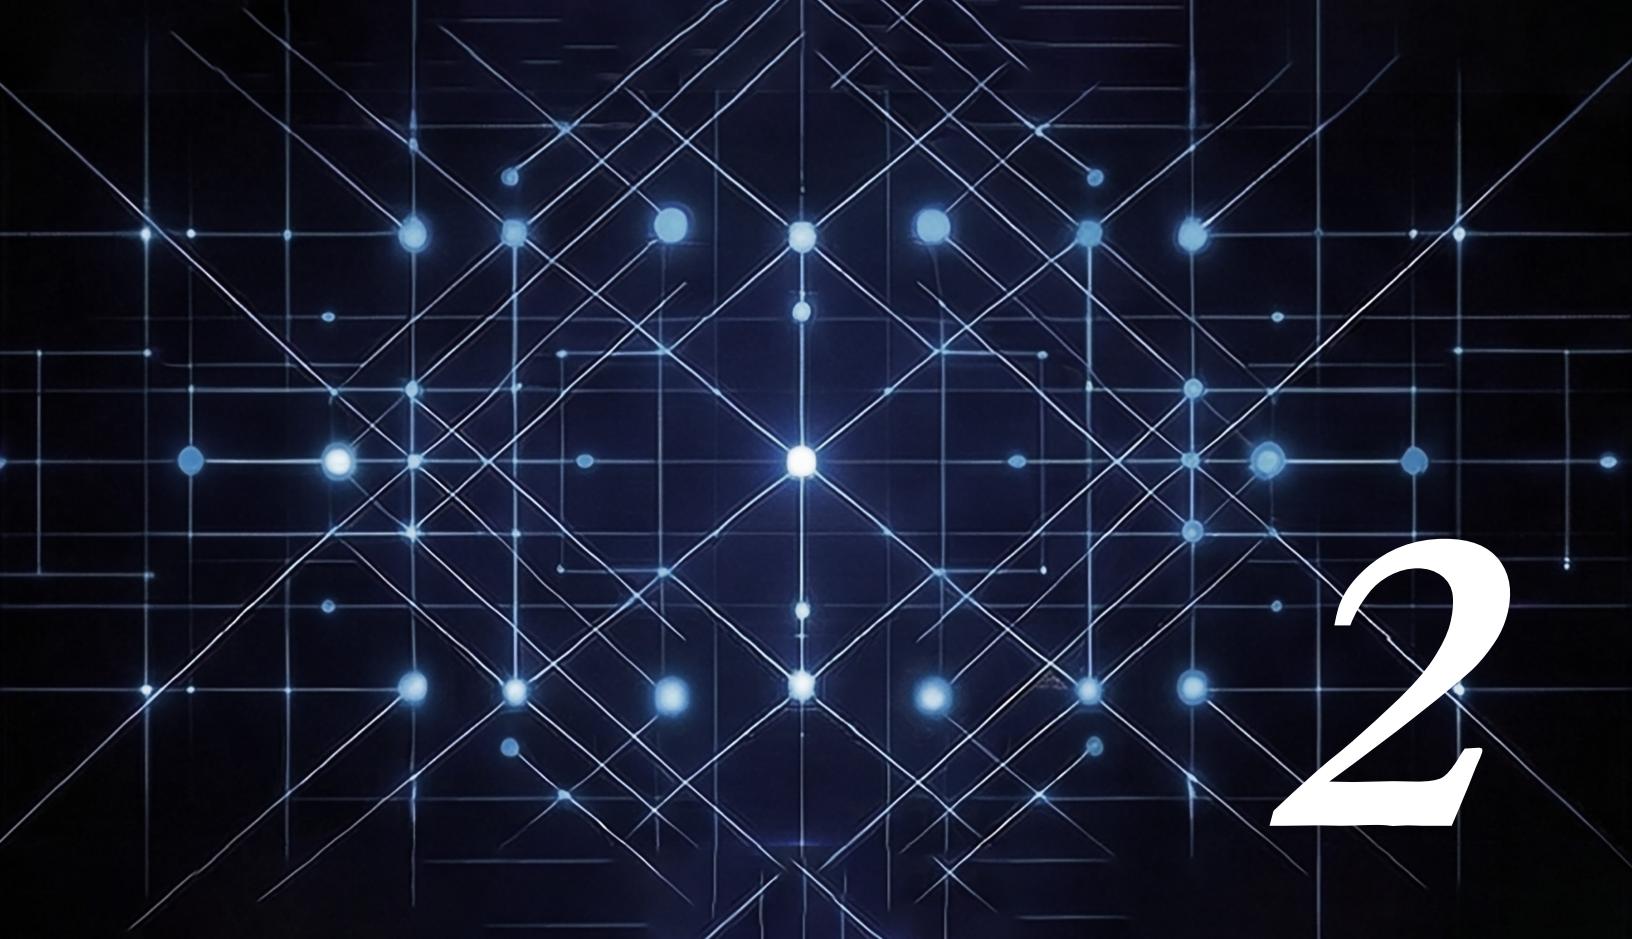
Insurers are backing away from AI coverage. In 2024 and 2025, major property and casualty carriers—AIG, WR Berkley, Great American—filed for regulatory approval to exclude AI liability from standard corporate policies.⁷ Why? Correlated failure risk, the absence of an actuarial baseline, and unpredictable failure modes are making default coverage untenable. In January 2026, Verisk—whose policy templates underpin 80% of the U.S property and casualty insurance market—introduced standard AI exclusion endorsements that will likely have the effect of making coverage gaps the default rather than the exception.⁸

As mainstream carriers back away from coverage for AI-related losses, some carriers are starting to offer affirmative AI coverage products, but only to organizations that can demonstrate ongoing robust AI governance practices. This suggests that the wave of exclusions represent a rejection of ungoverned AI adoption, not a categorical rejection of all AI use.

Insurers are aggressively deploying AI for their own operations. The irony is that the insurance industry outpaces nearly all other industries in enterprise AI adoption for their own internal operations: underwriting, claims processing, fraud detection, and customer service.⁹

This situation is best interpreted as a rational application of real information asym-

metry, not hypocrisy. Insurers deploying AI internally have visibility into their own governance processes. They know where AI authority stops and human judgment begins and have designed systems and processes accordingly. They can see their own risk. They can't see the risk of their customers. Until they have that visibility, they'd rather exclude than guess.



2

The Four Constraints & Five Delegation Categories

Key Points

- Four operational constraints—*independent of AI capability*—determine which AI systems can be deployed: Consequence of Error, Verification Cost, Accountability and Human Presence Requirements, and Physical Reality Limitations.
- When verification costs equal or exceed efficiency gains, the economic case for AI collapses—a phenomenon economists call “so-so automation” that displaces labor without generating net productivity gains.
- Certain tasks require human authorization or authentic presence that cannot be transferred to AI; human authenticity has measurable market value, with AI-generated content receiving 45% less engagement and 14% lower purchase consideration when detected.
- Five delegation categories define how humans and AI interact across tasks: Automated (AI with minimal oversight), Verified (human as editor), Co-Piloted (human engaged throughout), Assisted (AI prepares but humans decide), and Human-Only (safety-critical physical work or authentic presence required).

Four types of operational constraints determine which AI systems deploy and which ones don’t. These exist independently of AI capability—you can have a perfect model that still can’t be deployed because of operational requirements.

The Four Constraints

Consequence of Error. *What* are the costs *when AI is wrong*? A poorly worded social media post is likely benign. Payment processing errors are serious but reversible. Medical diagnostic errors can kill people. The level of consequence determines how much AI risk firms can tolerate.

Verification Cost. *Can someone double check AI’s output without redoing all the work?* Software code can be verified through affordable automated unit tests. Medical diagnoses require a highly paid physician to re-examine the case—potentially more expensive than doing the initial diagnosis themselves. When verification costs equal or exceed efficiencies gained, the economic case for using AI collapses. Economists call this “so-so automation”—technology that displaces labor without generating net productivity gains.¹⁰

Accountability and Presence. *Does the task require human authorization or authentic human presence?* A physician can’t transfer malpractice liability to an AI, even if the AI is more accurate. An engineer can’t delegate their Professional Engi-

neer (PE) stamp—the legal certification that a design is safe for the public. Other tasks require authentic presence—counseling, leadership, negotiation—where the human relationship is the point. This isn’t just sentiment: AI-generated content receives 45% less engagement than human-written equivalents, and purchase consideration drops 14% when consumers detect synthetic marketing.¹¹ Human authenticity has measurable market value.

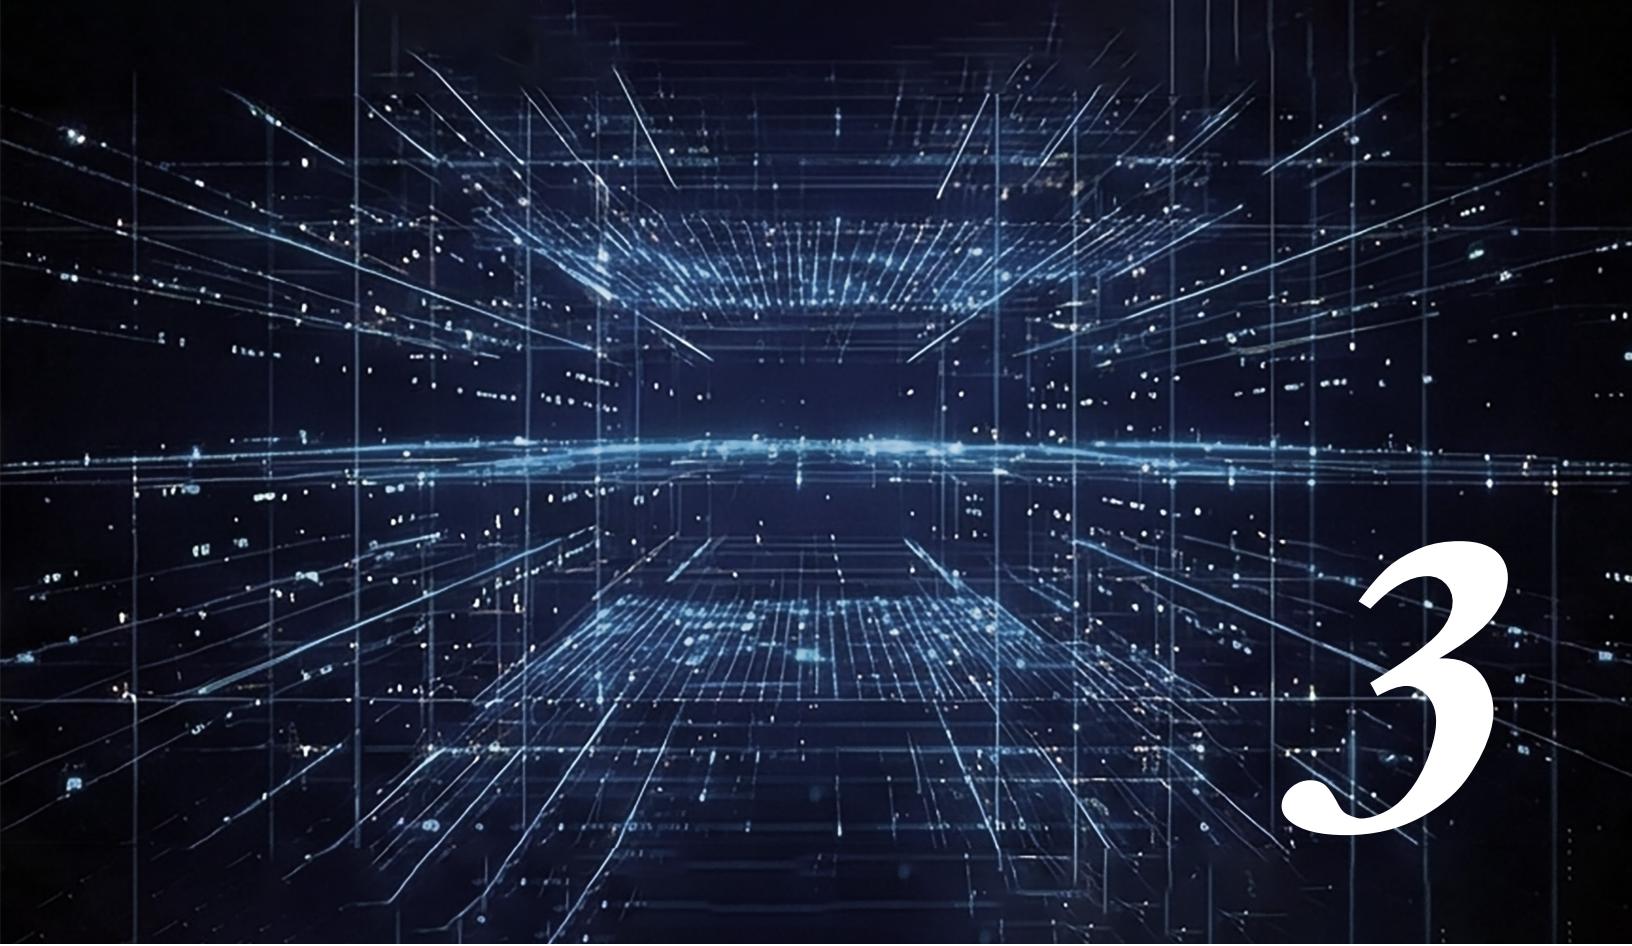
Physical Reality. *Does the task require manipulating atoms in unstructured environments?* Software scales infinitely, but physical intervention is constrained by safety and physics. The capital requirements for autonomous physical robotics systems are extraordinary—billions in R&D, sensor infrastructure, and safety validation. MIT research found only 23% of technically-exposed vision tasks were economically attractive to automate, given upfront costs.¹² Better models won’t close this gap soon. It’s a capital barrier.

The Five AI Delegation Categories

To see how the four operational constraints interact when applied to specific tasks that are part of real-world jobs, we developed five categories of AI delegation. Each category defines a state of human and AI interaction starting from fully AI-delegated automation tasks that require minimal human oversight to Human-Only tasks.

Delegation Category	Human Role	When It Applies
Automated	<i>Governor</i>	Low consequence, affordable verification, delegable, digital
Verified	<i>Editor</i>	Moderate consequence, deterministic checks, digital
Co-Piloted	<i>Pilot</i>	Expert verification is required; humans are engaged throughout
Assisted	<i>Judge</i>	High consequence or non-delegable; AI prepares, humans must decide
Human-Only	<i>Actor</i>	Safety-critical physical work or authentic human presence required

Automated and **Verified** work permits rapid delegation because verification is affordable or unnecessary. **Co-Piloted** and **Assisted** work improve productivity, but humans remain accountable. **Human-Only** work is fundamentally human.



3

The National Opportunity

Key Points

- Analysis of 18,898 tasks across 848 jobs held by 148 million U.S. workers estimates the total national opportunity for AI delegation at \$3.24 trillion annually (31.8% of U.S. base wages), with a sensitivity range of \$2.47–\$4.34 trillion.
- The governance-safe floor—work firms can safely delegate today with affordable, deterministic verification—is \$1.6 trillion annually, comprising \$1.02 trillion in core work and \$0.58 trillion in coordination efficiency gains.
- Co-Piloted workflows where AI assists but humans remain fully engaged yield \$0.52–\$1.74 trillion in potential gains, while Assisted workflows where AI accelerates preparation but humans make final decisions offer \$0.27–\$1.12 trillion.
- Human-Only work receives no direct efficiency value, but AI can still enhance quality through simulated practice environments and serving as a sparring partner for high-stakes scenarios like crisis leadership and difficult negotiations.

Starting from existing government databases, we developed an expanded and enriched task-level database covering approximately 18,898 tasks that define 848 jobs held by approximately 148 million U.S. workers. We sorted tasks by their centrality to the role, categorizing every activity as either **Core Work** (the specialized duties that define the role) or **Coordination Work** (the universal administrative overhead common to nearly all occupations). We then calculated the delegable share of wages after systematically assigning delegation categories to core tasks and applying a coordination coefficient to coordination tasks.

This process yielded a task level estimate of AI delegable wage mass based on realistic governance constraints. This process makes it possible to generate both quantitative and qualitative guidance for where and how organizations *can safely and effectively deploy AI* at the national, state, industry, firm, department, job, and task level.

\$3.24 Trillion Under Governance Constraints

We estimate the present total U.S. national opportunity for AI delegation to be **\$3.24 trillion annually (31.8% of U.S. base wages)** under base assumptions, with a sensitivity range of \$2.47 trillion –\$4.34 trillion.

NATIONAL AI OPPORTUNITY: EXPLODED VIEW

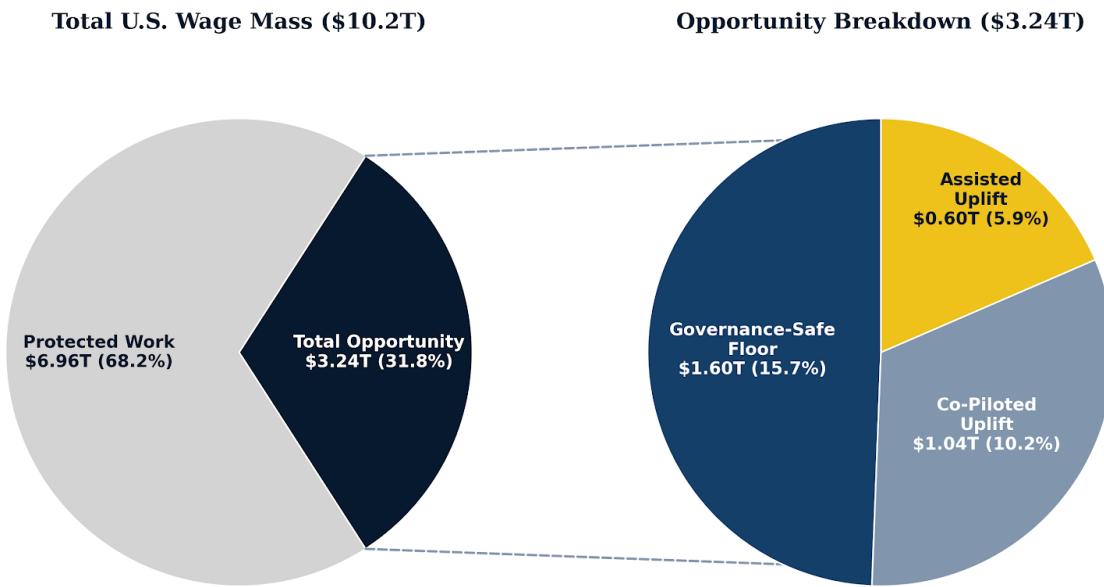


Figure 1 - National AI Opportunity in Context

The Governance-Safe Floor: \$1.60 Trillion Annually

We estimate the governance-safe floor for AI delegation in the U.S to be \$1.6 trillion annually. This floor represents work firms can safely delegate today because verification is affordable and deterministic:

\$1.02 trillion in core work delegation (10.0% of wages): Tasks for which checking the AI outputs costs far less than just doing the work—payment processing, data extraction, document classification, routine calculations.

\$0.58 trillion in coordination efficiency (5.7% of wages): Reducing the “work about work”—scheduling, status updates, email triage—that burdens every role. Research shows knowledge workers spend 57-60% of time on coordination rather than core work.¹³ AI fluency reduces this overhead even when core work remains human-performed.

This floor is conservative and immediately actionable.

Co-Piloted Uplift: \$0.52 Trillion to \$1.74 Trillion

The Co-Piloted Uplift represents productivity gains from human-anchored workflows where AI assists but humans remain fully engaged—code development, content creation, research. Controlled studies show 40-55% productivity gains in early AI co-pilot style deployments.¹⁴ The gains depend on the quality of verification processes and the sophistication of the redesigned workflow.

Conservative assumptions (15% net time reduction after verification overhead):
\$0.52 trillion

- Base assumptions (30% net reduction): \$1.04 trillion
- Aggressive assumptions (50% net reduction): \$1.74 trillion
- Assisted Uplift: \$0.27 Trillion to \$1.12 Trillion

Assisted Uplift: \$0.27 Trillion to \$1.12 Trillion

Gains where AI accelerates preparatory work—data extraction, option generation, analysis—but humans make final calls because accountability cannot transfer. This includes underwriting, diagnostic support, legal research, and strategic planning.

- Conservative assumptions (12% net time reduction): \$0.27 trillion
- Base assumptions (27% net reduction): \$0.60 trillion
- Aggressive assumptions (50% net reduction): \$1.12 trillion

Gains are smaller and more variable because verification often approaches the cost of re-derivation. Healthcare illustrates a particular risk: historical patterns of supplier-induced demand suggest that when AI reduces the cost of diagnostics, utilization may increase, creating downstream burden that offsets efficiency gains.¹⁵

Human-Only Work: Practice Environments & Sparring Partners

While we assign no direct efficiency value to Human-Only work, AI can still play a valuable role here. AI can provide simulated practice environments for humans to rehearse high-stakes scenarios—grief counseling, difficult negotiations, crisis leadership. It can also serve as a sparring partner, offering critiques and alternative perspectives that help humans sharpen their judgment. The value here isn't efficiency—it's quality improvement in the work where the human contribution matters most.



4

How Transformation Actually Happens

Key Points

- AI doesn't necessarily eliminate roles—it distills them, evaporating high-delegation tasks and concentrating human work on judgment, relationships, and accountability that only humans can provide.
- Role distillation requires active architectural planning; without deliberate reinvestment, freed human capacity is reabsorbed by coordination overhead and administrative drift rather than channeled toward high-value initiatives.
- Three illustrative workforce profiles reveal distinct AI opportunity patterns: financial services shows balanced opportunity across automation and advisory work; healthcare concentrates in Assisted Uplift due to non-transferable clinical liability; manufacturing sees Co-Piloted gains exceeding automation because physical work in unstructured environments remains protected while surrounding cognitive work benefits from AI augmentation.
- Each sector faces a strategic choice: banks must decide whether to deepen client relationships or let gains absorb into legacy processes; health systems must redesign care delivery or risk efficiency gains becoming unnecessary tests and documentation; manufacturers must invest in engineering AI fluency or miss augmentation opportunities by assuming they are “safe” from AI.

Distillation, Not Necessarily Elimination

To understand how governance constraints apply across sectors, we constructed three illustrative workforce profiles: a regional bank, a health system, and a custom equipment manufacturer. These are fictional organizations—not real companies—built using actual role structures, BLS wage benchmarks, and industry-standard workforce distributions. They demonstrate how the same data produces different deployment maps depending on workforce composition.

What the analysis reveals: *AI doesn't necessarily eliminate roles. It distills them.* High delegation potential tasks evaporate from daily workflows, allowing human work to concentrate on what only humans can do.

Capturing and Reinvesting Reclaimed Human Capital

However, role distillation is not a passive outcome; it is an active architectural choice. Without deliberate planning, the human capacity freed by AI evaporates—instantly reabsorbed by the ‘coordination tax’ of unnecessary meetings and administrative drift. Leaders must reverse the sequence: identify the high-value initiatives

that demand human ingenuity *first*, and then fund them with the capacity reclaimed from AI systems. You cannot bank efficiency; you must reinvest it.

The Teller Who Became an Advisor

ColdWest Bancorp is a fictional Midwest regional bank—\$25 billion in assets, 110 branches, 2,910 employees. Built on deep relationships with middle-market manufacturers and commercial real estate developers. A conservative credit culture that prizes relationship stickiness over geographic expansion.

Consider Maria, a composite of their 900-person teller workforce.



Ten years ago, Maria's day was transactions. Cash deposits. Check processing. Balance inquiries. Work that required accuracy and a pleasant demeanor, but not much judgment.

Today, ATMs and mobile banking handle 80% of those transactions. Maria's branch sees half the foot traffic it once did. But Maria still works there—and her role has transformed.

Now when customers come in, they come with problems. A small business owner whose cash flow doesn't match his loan covenants. A retiree confused by an estate transfer. A young couple who need someone to walk them through their first mortgage. Maria has become an advisor, a problem-solver, a relationship manager. The transactions evaporated. The judgment work concentrated.

This pattern—documented by economists studying automation’s historical effects¹⁶—is precisely what we mean by role distillation. The work didn’t disappear. It was purified.

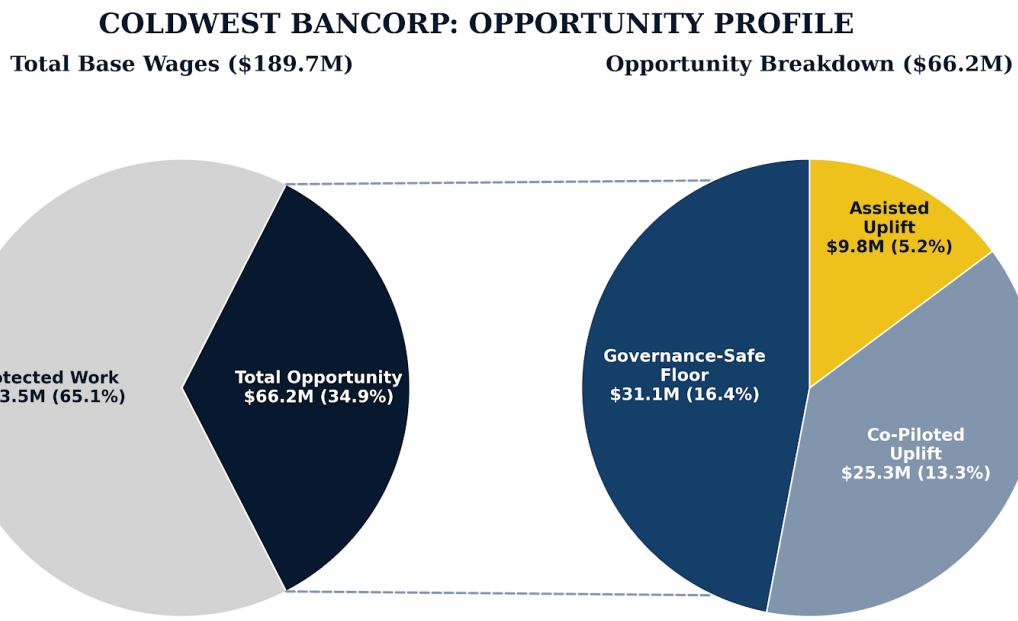


Figure 2 - ColdWest Bancorp AI Opportunity Profile

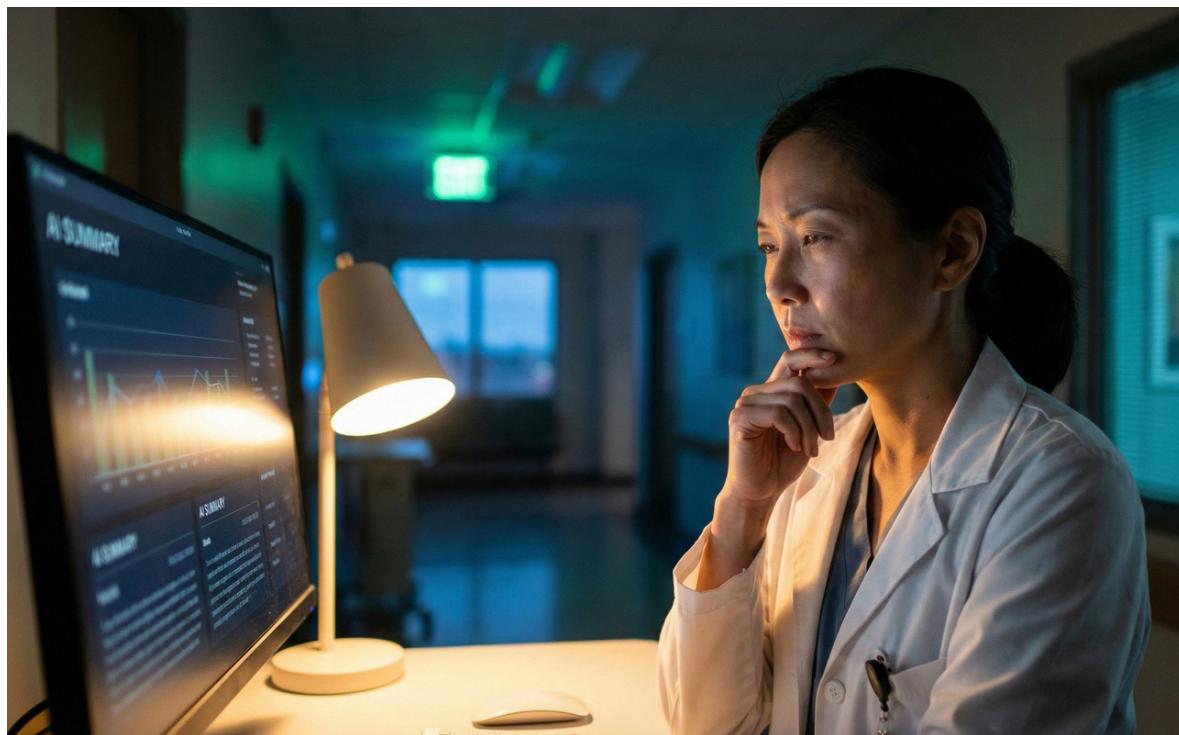
ColdWest shows balanced opportunity across all three tiers—substantial governance-safe automation, significant co-pilot gains, and meaningful decision support. This reflects the mix of routine transactions and judgment-intensive advisory work that characterizes financial services.

Our analysis found that the compliance function demonstrates a potential for non-linear scaling: Using AI to triage thousands of potential fraudulent transaction alerts helps financial institutions break free of the proportional relationship between customer volume and compliance headcount. The same analyst capacity can monitor much larger customer bases—eliminating the “compliance tax” that has historically constrained growth.

The question for ColdWest’s leadership: Will they use their new AI-enabled capacity to deepen client relationships and launch new advisory services, or will they allow these gains to be absorbed by legacy processes—missing the chance to shift their competitive value proposition?

The Physician Who Can't Delegate Liability

KaleidoHealth is a fictional non-profit health system in the Southeast—multiple hospitals, outpatient clinics, a 3,000-physician medical group. A clinical and economic anchor for its region, it has 32,900 employees and pays them \$2.58 billion in wages annually.



Consider Dr. Sarah Chen, an internist in KaleidoHealth's primary care network.

Dr. Chen spends 47% of her time on coordination—charting, inbox management, prior authorizations, referral letters, quality reporting. Nearly half her day is consumed by “work about work” before she sees a single patient.

But AI has started helping with much of this. Her Electronic Health Record (EHR) system now drafts visit summaries. An AI co-pilot helps with inbox triage. Research tools accelerate her literature reviews when patients present with unusual symptoms.

But here's what AI cannot do: sign her name to a diagnosis. Dr. Chen carries malpractice liability for every clinical decision. No algorithm, however accurate, can transfer that accountability. She can use AI to prepare—to synthesize research, generate differential diagnoses, surface relevant protocols—but the final call for each diagnosis remains hers. Legally. Professionally. Ethically.

This is why KaleidoHealth's opportunity profile looks so different from that of Coldwest Bancorp.

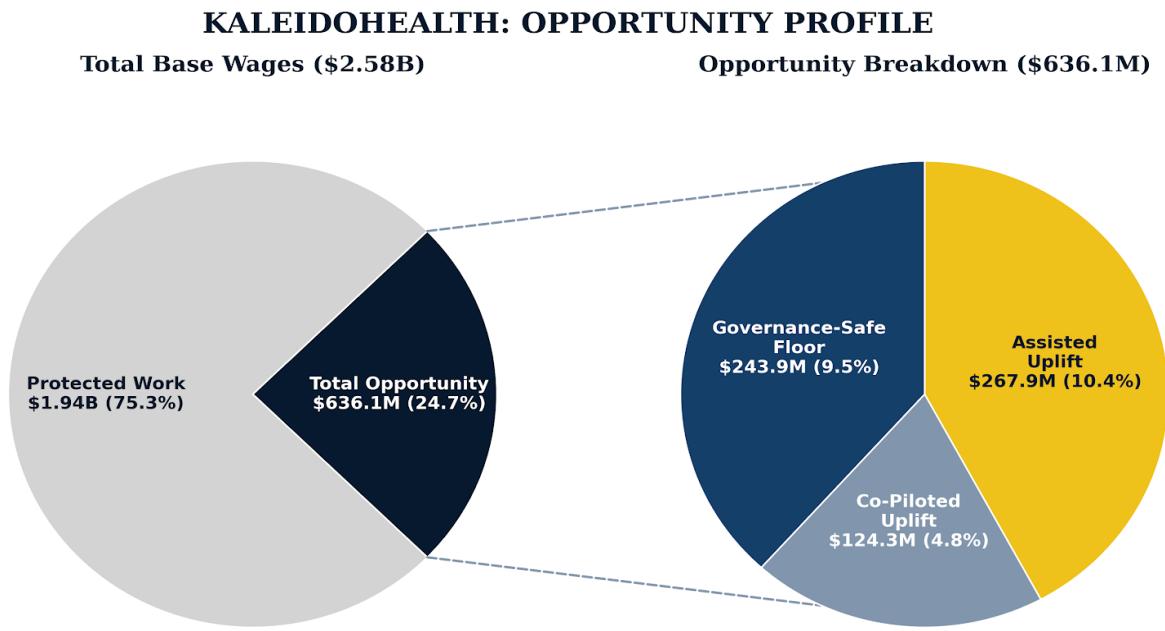


Figure 3 - KaleidoHealth AI Opportunity Profile

Assisted Uplift is the largest tier—10.4%. This is healthcare's unique signature. Clinicians make high-stakes decisions constantly, and AI can dramatically accelerate the preparatory work even when final judgment must remain human.

But there's a trap. Historical patterns of supplier-induced demand in healthcare suggest that when AI reduces the cost of diagnostics, utilization may increase—more tests ordered, more documentation generated.¹⁷ The efficiency gains can create downstream verification burden that offsets what was saved. Dr. Chen might spend less time writing notes but more time reviewing AI-generated drafts. The freed capacity doesn't automatically become patient time.

The Billing department shows a sharp contrast. Medical coders have high Automated and Verified task shares because their tasks follow structured rules verifiable against payer requirements. Here, role distillation is highly visible: AI handles routine coding, humans shift to exception handling, denial pattern analysis, and physician coaching on documentation practices—work that creates more value than mechanical code lookup.

The question for KaleidoHealth's leadership: Will they redesign care delivery—freeing clinicians to focus on complex, human-centered medicine—or will they let their efficiency gains become unnecessary tests and documentation that takes time away from patient connection?

The Welder Who Can't Be Replaced by Software

Pacific FoodPro is a fictional manufacturer in the Pacific Northwest—**1,361 employees, \$121.5 million in wages**. They build custom food processing equipment: conveyors, fillers, case packers, palletizers. Made-to-order systems for breweries, frozen-fruit packers, dairy processors. Each project requires close collaboration between engineers and food scientists to meet FDA and USDA sanitation standards.



Consider Marcus, a senior welder on the fabrication floor.

Marcus works with stainless steel—the material of choice for food-contact surfaces because it doesn't harbor bacteria. Every weld he makes must be smooth, continuous, free of crevices where microorganisms could hide. The specifications are exacting. The consequences of failure are recalls, contamination, lawsuits.

No robot is taking Marcus's job anytime soon. Not because robots can't weld—they can, and they do, in high-volume automotive plants where the same weld happens 10,000 times a day. But Pacific FoodPro doesn't make the same thing twice. Each system is custom. Each weld is different. Each day, Marcus encounters joints that don't quite match the drawings because the customer changed specifications mid-project, or because the steel supplier sent slightly different gauge material, or because the upstream fabrication team made an adjustment that engineering hasn't documented yet.

Marcus navigates this variability instinctively. A traditional robot would need to be

reprogrammed. The economics don't work. Current state-of-the-art AI robotics technology depends on highly structured environments and has such high cost and complexity that it is a poor fit for a custom manufacturer like Pacific FoodPro.

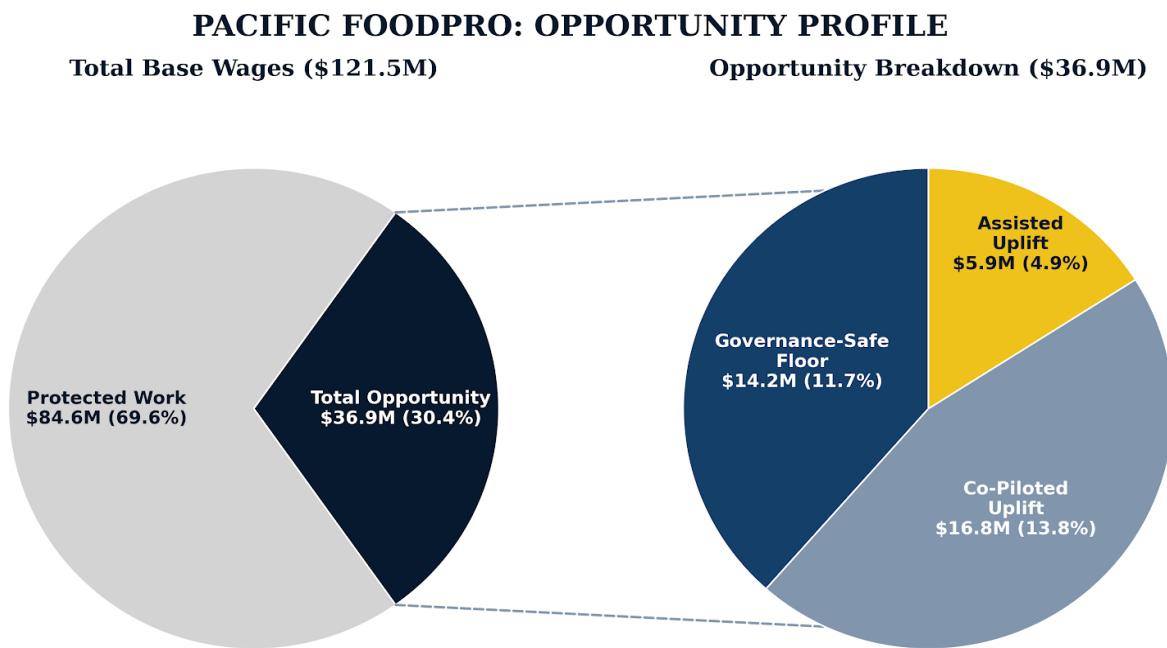


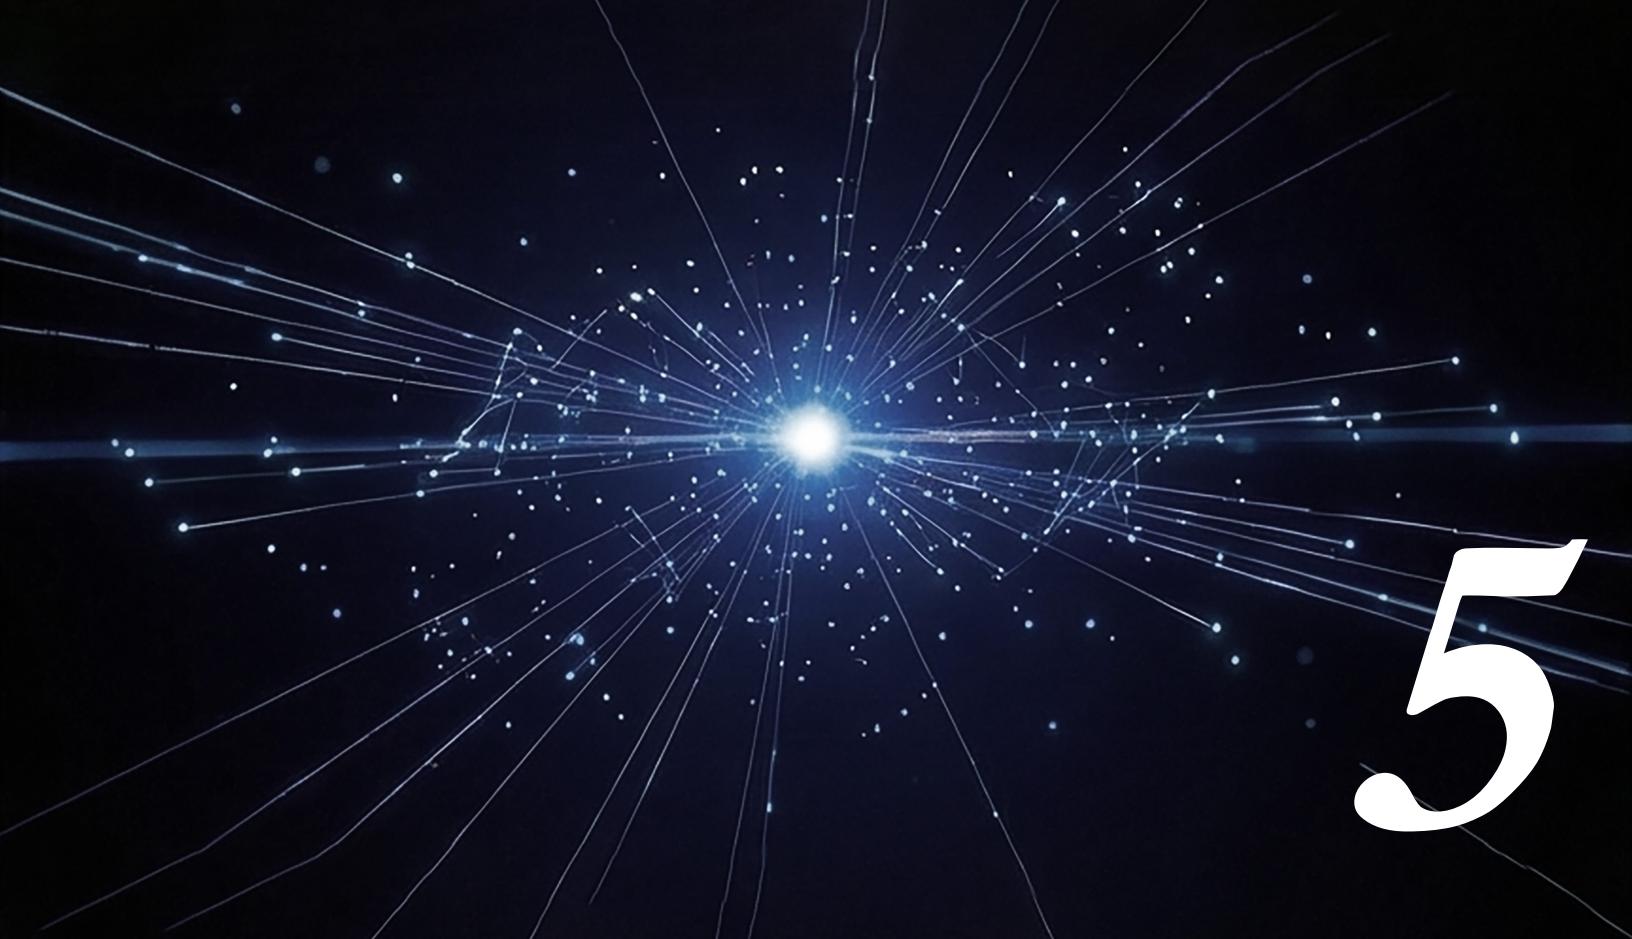
Figure 4 - Pacific FoodPro AI Opportunity Profile

Co-Piloted Uplift exceeds the Governance-Safe Floor—13.8% versus 11.7%. This is manufacturing's signature pattern. The physical work itself is protected by physical reality constraints in their unstructured environment. But the cognitive work surrounding it—design, planning, troubleshooting, documentation—benefits enormously from AI co-piloting.

Pacific FoodPro's engineers don't need AI to replace Marcus. They need AI to help them design systems faster, troubleshoot problems more efficiently, and generate documentation that used to take days or weeks to produce. The factory floor stays human, but the humans have better data faster to work from. The opportunity lies in augmenting the engineers and operators who manage those physical systems.

The question for Pacific FoodPro's leadership: Will they invest in AI fluency for their engineering and operations teams, or assume that manufacturing is “safe” from AI and miss the engineering augmentation opportunity entirely?

Based on our analysis of the national labor force and the constructed industry profiles, the potential efficiency gain from AI delegation is real and substantial. Where and how to safely and effectively introduce AI is increasingly clear.



5

Why Most Managers Will Struggle

Key Points

- Most managers will struggle with AI deployment because obvious approaches often fail: they target glamorous core work instead of the 57% coordination overhead where governance constraints are lightest, train for fear instead of fluency, and chase high-risk moonshots before building organizational confidence.
- The workforce splits into fearful non-users and overconfident slop-generators; a 2025 field experiment found AI reduced email processing time by 17%, but freed capacity was absorbed by organizational expansion rather than captured strategically—demonstrating that bad AI use can be worse than no AI use.
- High-performing organizations share distinct patterns: they establish reinvestment intent before deployment, fundamentally redesign workflows rather than just adding tools, concentrate 20%+ of digital budgets on AI, and understand the difference between technical exposure and safe deployment.
- Beyond efficiency, the ultimate prize is field-reshaping—using AI to change competitive rules across industries—but the path runs through operational fluency; firms that treat efficiency as the end goal will optimize their way to irrelevance.
- Historical patterns suggest AI will follow automation’s arc of net job creation through expansion, but outcomes depend on leadership intent: organizations can deliberately redesign work around human strengths and flourishing, or simply reshuffle which dehumanizing tasks humans perform.

It is also clear that most managers will likely struggle—not due to lack of effort, but because the obvious approaches often don’t work, and the correct but counterintuitive ones can feel wrong.

Failure Mode 1: Confusing Coordination with Core Work

Knowledge workers spend 57-60% of their time on coordination tasks—email, meetings, status updates, document search / production / management, scheduling.¹³ This is the low-hanging fruit. AI can meaningfully reduce this overhead without touching core work tasks at all.

But most managers don’t see it this way. They look at their marketing team and think “AI can write copy.” They look at their analysts and think “AI can build models.” They look at their engineers and think “AI can write code.” They go straight for core work—the visible, valued, identity-defining work that everyone cares most about.

This is backwards. Core work is where governance constraints bind hardest. Verifi-

cation is more expensive for core work. Consequences matter more. Accountability is less clear. Going straight for core work means going straight for the hardest problems in AI adoption.

Meanwhile, coordination work sits there, unglamorous and ignored: the 57% of everyone's day that nobody thinks about. The meeting scheduling, the status email drafting, the document search, the calendar management, the travel booking, the expense reporting. Work where errors are reversible, verification is affordable, and nobody's professional identity is threatened.

Managers who succeed start with coordination. They give people tools to reclaim the overhead tax before they touch the work people care about.

Failure Mode 2: Training for Fear Instead of Fluency

Most corporate AI training focuses on what not to do: Don't put confidential data in ChatGPT. Don't trust AI output without verification. Don't use AI for decisions that require human judgment. The result? The workforce splits into two camps.

Some workers, internalizing only the risks, become fearful non-users. They avoid AI altogether, uncertain about what is safe or valuable. They know the dangers but lack guidance on practical opportunities—warned about hallucinations but never taught to spot them, told to verify but never shown how, cautioned about appropriate use but never given role-specific examples.

Other workers treat the “don’ts” as a checklist for safe experimentation. Confident that avoiding forbidden actions equals fluency, they dive in—producing a flood of AI-generated content that is technically correct but contextually off-base. The result is an explosion of slop: emails that are longer, vaguer, and more generic than human-written ones; reports that miss the point; presentations that cover all the bases but make no argument. This output, while well-intentioned, creates new resentment among the fearful non-users who must now wade through the slop created by less cautious colleagues.¹⁸

The tension between these groups slows progress. True fluency requires more than knowing what to avoid; it demands understanding how to use AI thoughtfully, discerning when it adds value, and recognizing when human judgment is essential. Without this, organizations risk trading one set of inefficiencies for another.

A 2025 field experiment with 7,000 knowledge workers found that while AI reduced email processing time by 17%, time spent in meetings remained unchanged. The efficiency gains were absorbed by “offsetting responses”—organizational expansion that consumed the freed capacity before it could be captured strategically.¹⁹ This is

the Jevons Paradox applied to management: free up capacity without clear intent, and it gets consumed by the very expansion it enabled.²⁰

In fact, *this can be worse than no AI at all*, as content jams up every downstream process. Reviewers spend more time wading through AI bloat than they saved. Meeting prep takes longer because the pre-reads are longer. Decision-making slows because the analysis is thorough but not insightful.

Managers who succeed train for productive AI fluency. They teach people to use AI well, not just safely. They show what good looks like in specific workflows, measure output quality—not just volume—and understand that bad AI use is worse than no AI use.

Failure Mode 3: Chasing Moonshots Before Building Confidence

The business press is full of AI transformation stories. Massive cost savings. Breakthrough capabilities. Competitive moats. Industry transformation. The pressure to do something big is intense.

This leads managers to attempt high-risk, high-reward use cases first like AI-powered customer service that handles complex complaints, automated underwriting that makes lending decisions, or algorithmic hiring that screens candidates. These are cases where AI could make a big difference—and where failure is most visible, most consequential, and most likely to generate lawsuits and regulatory scrutiny.

These are exactly the wrong places to start. Some of them are excellent use cases—eventually. But they require organizational capabilities that don't exist yet. Verification processes that haven't been designed and built. Governance processes that haven't been tested. Cultural confidence that hasn't been earned.

Managers who succeed start simple. They build verification systems on low-stakes work. They earn organizational confidence through demonstrated success. They establish governance processes that can scale. Then—and only then—they move up the risk curve. By the time they attempt the moonshot use cases, they have the infrastructure to execute them.

The Minority Who Will Succeed

Analysis of AI high performers reveals what separates winners from the rest.

They have reinvestment intent before they deploy. 80% set growth or innovation objectives alongside efficiency goals. They know what they'll do with freed capacity

before they free it. Without a destination—market expansion, product development, relationship depth—efficiency gains evaporate into noise.²¹

They redesign workflows, not just tools. 55% fundamentally redesign workflows (versus 20% of others). They don't just add AI to existing processes. They rethink where the human-AI boundary should be, positioning AI to do what it does well while humans focus on what they do well.²²

They concentrate investment. They put 20%+ of their digital budget into AI (versus 7% of others). They're not dabbling. They're committing. And that commitment funds the verification architecture, training infrastructure, and governance systems that make deployment sustainable.²³

They understand the difference between exposure and deployment. Our analysis found 92% of U.S. wage mass has technical AI exposure—but only 15.7% is immediately safe to delegate. High performers don't confuse the ceiling with the floor. They start with what's safe, build capability, and expand systematically.

The total national AI deployment opportunity is approximately \$3.24 trillion. Most of it will be captured by a minority of organizations led by managers who understand these dynamics. The rest will announce AI initiatives, generate press releases, and wonder why nothing changed.

Beyond Efficiency: The Field-Reshaping Horizon

This paper focuses on operational reality—what can be safely deployed under current governance constraints. But efficiency gains aren't the ultimate prize. The ultimate prize is field reshaping: using AI to change the rules of competition itself, a concept platform strategist Sangeet Paul Choudary introduced in his book *Reshuffle: Who wins when AI restacks the knowledge economy.*²⁴

Choudary convincingly argues that the biggest opportunity isn't optimizing existing workflows—it's recognizing how AI changes coordination mechanisms and control points across entire industries. Companies like Uber Freight restructure trucking by algorithmically controlling pricing, routing, and load assignments. They don't just compete better; they reshape the competitive landscape so others must adapt to rules they've established.

Unfortunately, most organizations aren't yet equipped to reshape their industries through AI. They lack the AI fluency, operating leverage, and organizational con-

fidence to make such transformational bets. But that doesn't mean they should ignore the field-reshaping horizon.

The path to field-reshaping runs through operational fluency. Efficiency gains from the governance-safe floor today fund more ambitious transformation tomorrow. Experience with delegation boundaries, verification systems, and human-AI collaboration builds the intuition that field reshaping requires. Firms don't become a field reshaper by announcing a transformation initiative. They become one by building the operational capability that makes transformation executable.

Leaders who capture the governance-safe opportunity while keeping field-reshaping possibilities in view will be positioned to act when their fluency and resources allow. Those who treat efficiency as the end goal will optimize their way to irrelevance.

The Employment Arc: Distillation Creates Jobs

We assert that the default narrative—AI eliminates jobs—misreads the historical pattern. Automation has historically enabled economic expansion that leads to net-new job creation.

ATMs didn't reduce teller employment. They reduced branch operating costs, enabling banks to open more branches, which employed more tellers doing different work.²⁵ Spreadsheets didn't eliminate accountants. They enabled more sophisticated analysis, creating demand for more accountants doing higher-value work. The pattern repeats: automation of routine tasks enables expansion that creates net new employment.

Early evidence suggests AI follows the same pattern. Organizations with a growth mindset are positioned to use freed capacity to fund expansion, which often leads to creation of new, higher-value jobs. With careful planning and execution, these new jobs can be more human than what AI displaced—relying on judgment, relationships, creativity, accountability that only humans can provide. These aren't consolation jobs. They're better jobs.

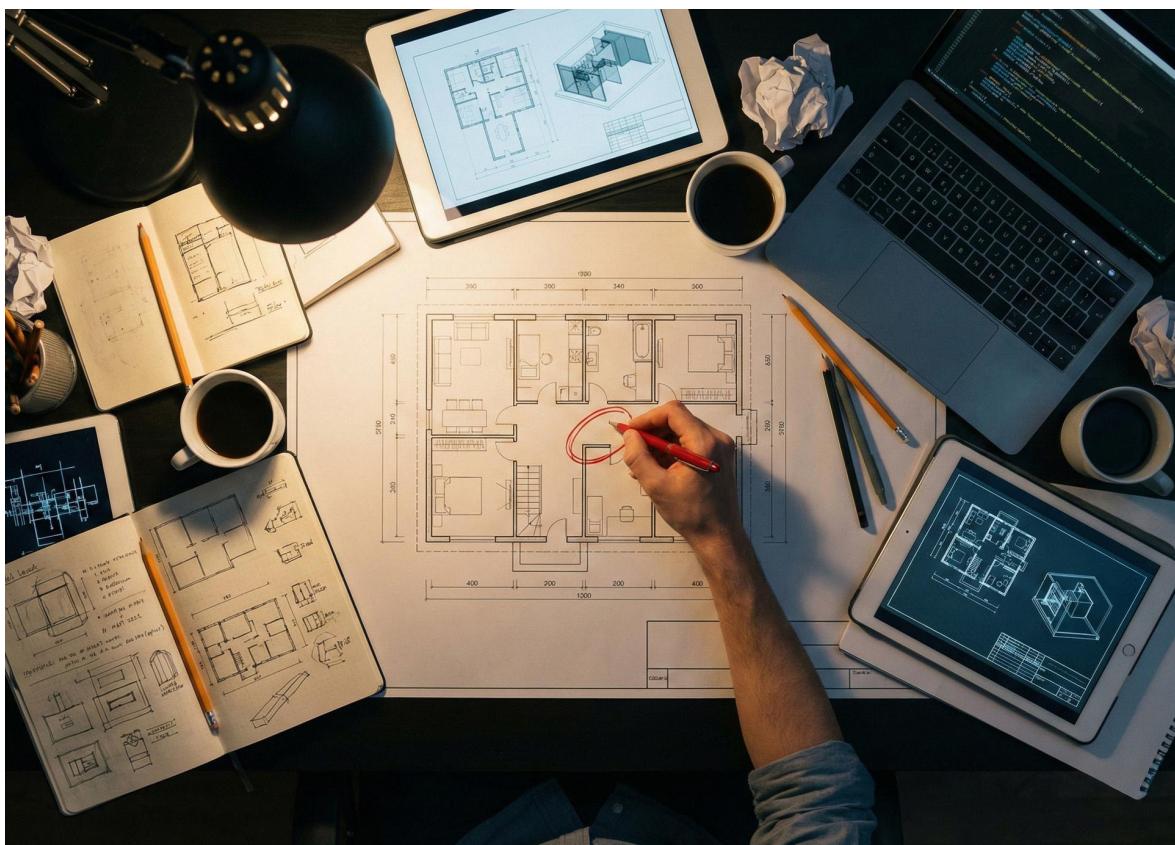
This doesn't happen automatically. Organizations that treat AI as a cost-cutting tool will cut costs—and headcount. But organizations with growth intent use freed capacity to fund expansion: new markets, new products, deeper customer relationships. Expansion requires people. More people, doing more human work, than before.

The job creation won't be evenly distributed. It will concentrate in organizations led by managers who understand that efficiency is fuel for growth, not an end in itself. The policy question isn't how to slow AI adoption to preserve existing jobs. It's how to accelerate the transition from distillation to growth so that new jobs emerge faster than old tasks evaporate.

The Opportunity to Make Work More Human

Here's a truth that the efficiency conversation obscures: many modern jobs are collections of tasks that ended up with humans not because humans find them meaningful or are particularly good at them, but because they couldn't be automated by the machines of the industrial and digital eras. They ended up with humans by default.

Consider TSA x-ray operators, rotated every 20 minutes because humans can't sustain the machine-like attention the task requires. Consider data entry clerks, call center agents reading scripts, workers monitoring dashboards for anomalies. These jobs don't play to human strengths—judgment, empathy, creativity, meaning-making. They play to human weaknesses, and we build elaborate structures to compensate for our limitations.



We believe AI creates an opportunity to change this. As high delegation potential tasks migrate to machines, what remains with humans could be deliberately designed around human strengths. Not the leftover work that AI can't do, but work that humans are uniquely suited to do well—and that they find genuinely

meaningful.

This won't happen automatically. Left to inertia, organizations will simply assign humans whatever tasks AI can't handle, recreating the same accidental job design that produced dehumanizing work in the first place. But with intention, leaders can use this moment to redesign work around human flourishing: roles built on judgment, relationships, creativity, and accountability rather than sustained attention, error-free repetition, and compliance with procedures.

The default narrative frames AI as a threat to human dignity—automation as alienation, efficiency as displacement. But the role distillation we document herein points toward a different possibility: work that is more human, not less. Roles that play to our irreducible strengths rather than compensating for our limitations. Jobs worth doing, not just jobs that remain.

This is the choice that AI refactoring of work presents: whether to put in the extra effort to make work more human or simply reshuffle which dehumanizing tasks humans perform. The technology doesn't decide. We do.



6

Implications for Policy Makers

Key Points

- State economic development officials should prepare workers for role distillation rather than displacement, funding programs that help workers transition to transformed roles—medical coders becoming physician coaches, processors becoming advisors, operators becoming system managers.
- Economic development incentives should be conditioned on reinvestment and growth plans, not just efficiency gains; subsidizing cost reduction without growth requirements funds headcount elimination without economic benefit.
- Insurers are effectively privatizing AI regulation by making governance a prerequisite for coverage, converting abstract “responsible AI” principles into technical requirements for insurability—creating both opportunity and risk for state regulators.
- Insurance regulators should clarify liability frameworks across delegation categories: confirm that Automated and Verified work fits existing frameworks, define what “active engagement” means for liability in Co-Piloted contexts, and set explicit high bars for autonomous systems in unstructured human environments.

For State Economic Development Officials

Workforce composition determines exposure. States with heavy manufacturing and healthcare concentrations have large “protected cores”—physical work, non-delegable clinical authority. These jobs aren’t disappearing; they’re distilling.

Prepare workers for role distillation. Policy should fund programs that prepare workers for transformed roles—medical coders becoming physician coaches, processors becoming advisors, operators becoming system managers. Workers aren’t moving from employment to unemployment—they’re moving from one version of a role to another.

Condition incentives on growth, not just efficiency. Economic development credits should be conditional on reinvestment plans, not just efficiency gains. Align incentives with growth initiatives, and AI adoption creates competitive advantage. Align them only with cost reduction, and you’ve subsidized headcount elimination without economic benefit. The historical pattern—automation enabling growth that creates net new employment—only holds when organizations reinvest freed capacity rather than pocket it.

Track job creation, not just job preservation. The policy conversation focuses on

jobs at risk. It should also track jobs being created—and whether those new roles play to human strengths or simply assign humans whatever AI can't do. We believe states that attract employers committed to human-centered job design will build more resilient, higher-quality workforces than those that compete only on cost.

Identify field-reshaping opportunities. Some industries are ripe for AI-enabled coordination changes that could create regional competitive advantage. States with strong logistics infrastructure might attract companies reshaping supply chain coordination. States with healthcare concentrations might become hubs for AI-assisted clinical workflow innovation. The question isn't just which jobs are protected—it's which transformations your state is positioned to lead.

State-specific AI delegation analysis—examining local workforce composition against the four constraints—can reveal which sectors face pressure, which are structurally protected, and where transformation opportunity concentrates.

For State Insurance Regulators

Insurers are making governance a prerequisite for coverage. To access affirmative AI insurance, organizations increasingly must deploy specific governance platforms providing real-time monitoring to underwriters. Insurers are effectively privatizing AI regulation—converting abstract principles of “responsible AI” into technical requirements for insurability.

This creates both opportunity and risk.

Clarify what's insurable. Automated and Verified work—purely digital, low-consequence, or human-verified—fits within existing liability frameworks. Removing ambiguity accelerates safe adoption.

Define governance requirements for lower-delegation potential work. Establish what “active engagement” (Co-Piloted) and “assisted preparation” (Assisted) mean for liability. If a physician uses AI for differential diagnosis, clarify that liability remains with the physician, provided AI was used as a tool, not a decision-maker.

Set explicit requirements for Human-Only contexts. The bar should be high for autonomous systems in unstructured human environments. Define testing, geofencing, and insurance thresholds required for deployment.

Account for compliance costs. EU AI Act compliance costs approximately €29,000 annually per high-risk AI system, with certification adding €17,000-23,000.²⁶ Governance isn't free. Premiums and reasonable-care standards should reflect this reality.

Mandate AI governance training. Require that licensed professionals—physicians, engineers, brokers—maintain continuing education on AI failure modes and verification accountability. You cannot govern what you do not understand.

About This Research

Building from the BLS ONET task and occupation database, we developed an expanded and enriched task-level database covering 18,898 tasks that define 848 jobs held by approximately 148 million U.S. workers. We sorted tasks by their centrality to the role and systematically assigned delegation categories.

The assessment protocol used a multi-model LLM council (Gemini, ChatGPT, Claude, Llama) to control for single-model bias, achieving Fleiss' Kappa 0.81 (AI system classification) and Cronbach's Alpha 0.88 (delegation potential scoring)—meeting or exceeding standard thresholds for social science measurement instruments.²⁷

We then used BLS OEW data to calculate the share of wages that corresponds to each task. This process yielded a task level estimate of AI delegable wage mass based on realistic governance constraints.

National numbers represent typical organizations, not best-in-class. Co-Piloted and Assisted layers are reported as scenario ranges because actual results depend on governance maturity and risk tolerance.

The methodology scales from firm-level AI delegation opportunity mapping to state-level economic analysis to national labor market modeling. We applied it to U.S. data; the approach extends to any economy where occupational task data is available.

We welcome collaboration with partners committed to joint research and implementation. Our enriched ONET database—848 occupations and 18,898 tasks combined with BLS OEW data—along with complete technical methodology, is available to engaged partners. Send inquiries to research@seampoint.com.

Suggested Citation: Seampoint LLC. (2026). *The Distillation of Work: Where AI Opportunity Concentrates and How Leaders Capture It*.

This research was conducted independently and funded by Seampoint. The Seampoint management team has no concentrated financial interest in any AI vendor or any firm mentioned in this report.

About Seampoint

Seampoint is a research and advisory firm built on the insight that value concentrates at the seam points—the critical boundaries where human authority meets AI capability. Properly designed seams help firms gain maximum AI advantage while appropriately managing risk.

We help firms build lasting advantage by identifying where AI delegation is appropriate, suggesting proven design patterns for utilizing specific AI technologies, designing governance protocols for safe delegation, and building the workforce fluency required to achieve breakthrough performance on both sides of the seam.

seampoint.com



Endnotes

- 1 Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). "GPTs are GPTs: An early look at the labor market impact potential of large language models." *arXiv*.
- 2 McKinsey Global Institute. (2023). "The economic potential of generative AI." McKinsey & Company.
- 3 U.S. Census Bureau Business Trends and Outlook Survey (BTOS), December 2025. Survey question: "Does this business currently use AI in any of its business functions?"
- 4 Financial Times (2024), "Insurers retreat from AI cover as risks mount"; Insurance Journal (2024); WR Berkley Corporation (2025), "Absolute AI exclusion filing."
- 5 BBC News. (2024). "Air Canada chatbot promised a discount it couldn't deliver."
- 6 Seyfarth Shaw LLP. (2024). "Mobley v. Workday: AI hiring discrimination case."
- 7 Financial Times. See note 4.
- 8 Verisk/ISO. (2025). Generative AI Exclusion endorsement (form CG 40 47), effective January 2026. Templates underpin ~80% of the U.S. property and casualty market.
- 9 BCG. (2025). "Insurance Leads in AI Adoption." Insurance outpaces nearly all industries in GenAI adoption.
- 10 Acemoglu, D. & Restrepo, P. (2019). "Automation and new tasks: How technology displaces and reinstates labor." *Journal of Economic Perspectives*. Introduces "so-so automation" framework.
- 11 Originality.ai (2025). "Over 1/2 of Long Posts on LinkedIn are Likely AI-Generated Since ChatGPT Launched." October 28, 2025. Raptive (formerly AdThrive) (2025). "AI Trust Study." Survey of 3,000 U.S. adults found 14% drop in purchase consideration and 14% decrease in willingness to pay a premium for products associated with AI-generated marketing content, along with 50% drop in content trust when AI authorship is suspected.
- 12 Svanberg, M., et al. (2024). "Beyond AI exposure: Which tasks are cost-effective to automate with computer vision?" MIT. Only 23% of AI-exposed vision tasks are economically attractive given upfront costs.
- 13 Microsoft. (2024). "Work Trend Index Annual Report." (60% on email, chat, meetings); Asana. (2024). "Anatomy of Work Global Index." (58% on "work about work").
- 14 Noy, S. & Zhang, W. (2023). "Experimental evidence on the productivity effects of generative artificial intelligence." *Science*. (40% faster); Peng, S., et al. (2023). "The impact of AI on developer productivity: Evidence from GitHub Copilot." *arXiv*. (55% faster).
- 15 Dhanoa, D., et al. (2013). "The evolving role of the radiologist." Canadian Association of Radiologists Journal. Vancouver General Hospital: 98% CT productivity gain, yet 60% utilization increase (2000-2008). Rao, V. M., et al. (2011). Medicare Part B utilization data. *Frontiers* (2025). "Studying supplier-induced demand in healthcare: an econophysics approach." *Frontiers in Health Services*. MDPI (2023). "Payment Systems, Supplier-Induced Demand, and Service Quality in Credence Goods: Results from a Laboratory Experiment." *Games*. Applied to AI: theoretical application suggests efficiency gains may be offset by downstream verification burden.

16 Bessen, J. (2015). *Learning by Doing: The Real Connection Between Innovation, Wages, and Wealth*. Yale University Press; Autor, D. (2024). “Applying AI to rebuild middle class jobs.” NBER Working Paper No. 32140.

17 Dhanoa et al. See note 15.

18 Bernstein, E., & Gerpott, F. H. (2025, September 24). AI-generated “workslop” is destroying productivity. *Harvard Business Review*.

19 Dillon, E., Jaffe, S., Immorlica, N., & Stanton, C. (2025). “Shifting work patterns with generative AI.” NBER/Microsoft. Randomized controlled trial, 7,137 workers across 66 firms.

20 Jevons, W. S. (1865). *The Coal Question: An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal-Mines*. Macmillan.

21 McKinsey & Company. (2025). “The state of AI in 2025: Agents, innovation, and transformation.”

22 Ibid.

23 Ibid.

24 Choudary, Sangeet Paul. (2025). *Reshuffle: Who wins when AI restacks the knowledge economy*. Amazon Digital Services LLC – KDP.

25 Bessen and Autor. See note 16.

26 Bignami, E. G., et al. (2025). “Balancing Innovation and Control: The European Union AI Act in an Era of Global Uncertainty.” *JMIR AI* 4:e75527. DOI: 10.2196/75527. Cites Renda, A., et al. (2021). “Study to Support an Impact Assessment of Regulatory Requirements for Artificial Intelligence in Europe.” Publications Office of the European Union. High-risk AI system: €29,277 (US \$34,153) annually; certification: €16,800-23,000 (US \$19,598-26,831) per unit.

27 Landis, J. R. & Koch, G. G. (1977). “The measurement of observer agreement for categorical data.” *Biometrics*. Kappa 0.81+ = “almost perfect agreement.”

Seampoint
January 2026
© 2026 Seampoint LLC
All rights reserved.

seampoint.com