

June 2026

The Future of Work: A Framework for Action Under Uncertainty

Mapping the AI workforce transition:
scenarios, narratives, policy gaps,
and paths forward

An Editorial Note from Fathom

AI's potential impact on the labor market has become one of the most urgent—and most contested—questions in the public debate about the technology. The stakes are high: work is how most people structure their days, develop competence, connect with others, and answer the question of who they are.

The public concern is real and measurable. A [2026 Quinnipiac University poll](#) found that 70% of Americans think AI advancements will lead to a decrease in job opportunities; among employed Americans, 30% worry their own jobs will become obsolete. Fathom's own polling found that nearly half of Americans agree that AI will disrupt the labor market more than previous technologies—and the share who believe it will create new opportunities has fallen, from 43% in 2024 to 37% in 2025.

Policymakers are responding: of the roughly 1,200 AI bills introduced across all 50 states in 2025, 145 became law, and employment is the only area where a significant share has done so.

The stakes extend beyond economics. Pope Leo XIV's encyclical *Magnifica Humanitas*, published in May 2026, insists that work “expresses and enhances the dignity of our lives,” and AI-driven displacement is not just an economic problem but an attack on human dignity.

The question is not only whether jobs will exist, but what workers are owed—and what human lives should look like in an economy shaped by AI.

Three questions sit at the center of the debate: how—and at what pace and scale—will AI affect the labor market? Should we do anything about it? And if so, what can we do to prepare?

We are not attempting to answer the first question. No one can, at least not with complete certainty. We approach the question of the future of work with humility: the evidence is fragmentary and the technology is moving faster than measurement. What we can say is that the range of plausible futures is wide, the policy responses they demand are very different, and the institutions required to act, in any scenario, take years to build. And humility is not the same as paralysis: the case for acting now is not about betting on any one future. It is about building the foundations that every future will need. We expect 2027 and 2028 to begin revealing which scenario we are in, and some of the responses it demands could already be years behind schedule.

This report maps the landscape: the data we have and what we're missing, four scenarios for how AI might reshape work, six narratives about what to do, the policy toolkit and its gaps, and the cross-cutting priorities that hold up across futures.

What History Can Teach Us—and Why It May Not Be Enough

Historical Precedents

AI shares important features of earlier technologies: it is disruptive, general-purpose, and has the potential to raise productivity across many sectors. Four historical precedents are often invoked in workforce disruption analysis, each carrying a governance lesson:

1. Industrialization (late 19th–early 20th century)

Displaced manual and artisanal labor while creating factory, service, and later professional work. Adjustment depended on the gradual expansion of mass education, labor protections, and social insurance—and where those institutional responses lagged, the result was social unrest.

↳ **The lesson:** Technological transitions require new institutional infrastructure that markets don't produce on their own.

2. The Great Depression

While not a technology story *per se*, it remains the most influential precedent for public action when private demand fails. Public works and conservation programs demonstrated that governments can credibly intervene at scale when private demand collapses.

↳ **The lesson:** Government can serve as an employer of last resort when markets fail, but only if the institutional capacity is built before the crisis.

3. Late-20th-century deindustrialization and the China shock

Saw manufacturing decline, with job loss outpacing job creation in import-competing regions of the U.S. Midwest and the South, depressing wages and labor-force participation for years. The federal response, Trade Adjustment Assistance, was underfunded, hard to access, and produced weak outcomes.

↳ **The lesson:** Skills retraining alone isn't effective unless it's sufficiently funded and geographically targeted.

4. Recent automation in services and routine office work

This has been linked to labor-market polarization: job growth at the top and bottom, pressure on the middle. Skills-based retraining helped some workers move up but did not absorb the losses, particularly for older workers and those with less formal education.

↳ **The lesson:** Skills programs help some workers, but certain groups predictably lag unless those programs are paired with income support.

What History Can Teach Us—and Why It May Not Be Enough

Why AI may be different

Across all four, the pattern is consistent: the institutional response lagged the disruption, retraining alone was insufficient, and which workers would fall through the cracks—those with the least institutional support, the fewest transferable skills, and the least mobility—was broadly predictable in advance.

These precedents are instructive, but AI may differ from previous disruptions in several ways:

Speed

Earlier tech-driven disruption waves played out over decades, giving institutions and workers time to adjust. AI's deployment cycle is years, even months, not generations, thus compressing the timeline adjustment depends on.

Cognitive breadth

Past automation primarily targeted routine manual work, and new jobs emerged in services and knowledge work. AI is general-purpose, and capable of following workers into those new “knowledge” jobs.

Intrinsic cognitive ability

As capabilities improve, displacement may extend along the dimension of cognitive ability itself, closing off the historical pathway by which workers moved up when manual labor was automated.

Constant improvement

Roles that look safe today may become vulnerable as capabilities expand. The bottlenecks where it once looked like humans and AI would remain permanently complementary can close as the models' abilities expand quickly.

The eroding tax base

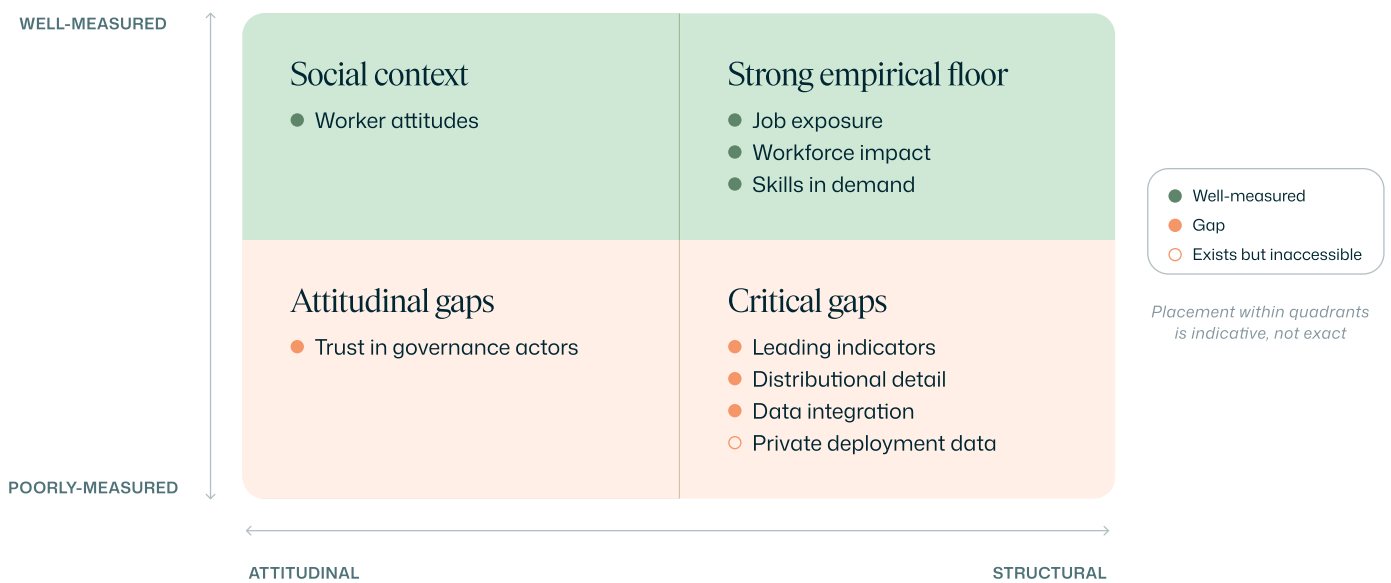
If AI shifts more income from labor to capital, payroll- and labor-income-based funding for the social safety net erodes precisely as demand on that safety net grows.

These differences do not guarantee unprecedented disruption. AI may yet follow the pattern of past general-purpose technologies, with significant but manageable restructuring over time. But the optimistic case—the “let the market sort it out” narrative, which we explore later in this report—rests on conditions earlier transitions enjoyed: time for markets and workers to adjust, durable niches where human work remained essential, and wages that tracked productivity. These are precisely the conditions AI's distinctive features may erode, and that is what makes the risks of inaction unusually asymmetric.

What we know and what we're missing

A central challenge in AI workforce policy is that the data is fragmented and the technology is moving faster than measurement is advancing. The chart below maps data sources along two axes: **how well each area is currently measured**, and **whether it captures attitudinal signals** (what workers think and feel) or **structural signals** (what's happening in the labor market).

The Evidence Landscape



Areas that are relatively well-measured

(as of mid-2026)

Job exposure and impact: Several ongoing sources track AI's deployment across tasks, occupations, and firms. They show a workforce widely exposed on paper but unevenly affected in practice: adoption is far higher among individuals (54.6% of working-age U.S. adults) than firms (18%). So far, observed use to date leans toward assisting workers more than replacing them.

Skills in demand: Well-tracked through job-posting and labor-market data, which show demand for AI skills growing fast—postings citing AI literacy up more than 70% year-on-year—but concentrated at the higher-paid end and still narrow overall, appearing in just 2.5% of U.S. job postings.

Worker attitudes: Extensively surveyed across multiple pollsters, with a consistent finding: use and anxiety are rising together. Workers report more worry than hope about AI's workplace impact, and trust in both companies and public institutions to manage the transition fairly remains low.

Gaps: where the data is missing

Leading indicators: There are no agreed indicators on AI-related layoffs, entry-level hiring shifts, or wage suppression that would signal the need for early action (support, retraining, or revision of deployments) before damage accumulates. Government statistics are built around lagging quarterly cycles.

Distributional detail: Impacts are typically reported by occupation, sector, or education level, but not often combined with demographic and geographic breakdowns (such as older workers in specific regions), and where it exists, it tends to come from academic research rather than official statistics. We need more of the constituency-level granularity that would move policymakers to act.

Data integration: Exposure data tells us where AI could be used; usage data and employer surveys tell us where it is being used. These layers are rarely integrated, especially at the national scale.

Private deployment data: The most reliable task-level and usage data lives inside private systems, primarily frontier labs and large employers. No standard frameworks exist for sharing it safely with researchers or governments. The actors deploying AI hold the clearest picture of its effects, while regulators, unions, and the public see only a partial one.

Trust in governance actors: Public sentiment research finds broad skepticism of the government's ability to regulate AI, but who the public would trust specifically to lead an AI-workforce transition—employers, federal agencies, state governments, unions, or independent bodies—remains an open question. Understanding this dimension of public sentiment is key, since any policy package depends on a legitimate convening actor, and it is not yet clear which actors carry that legitimacy.

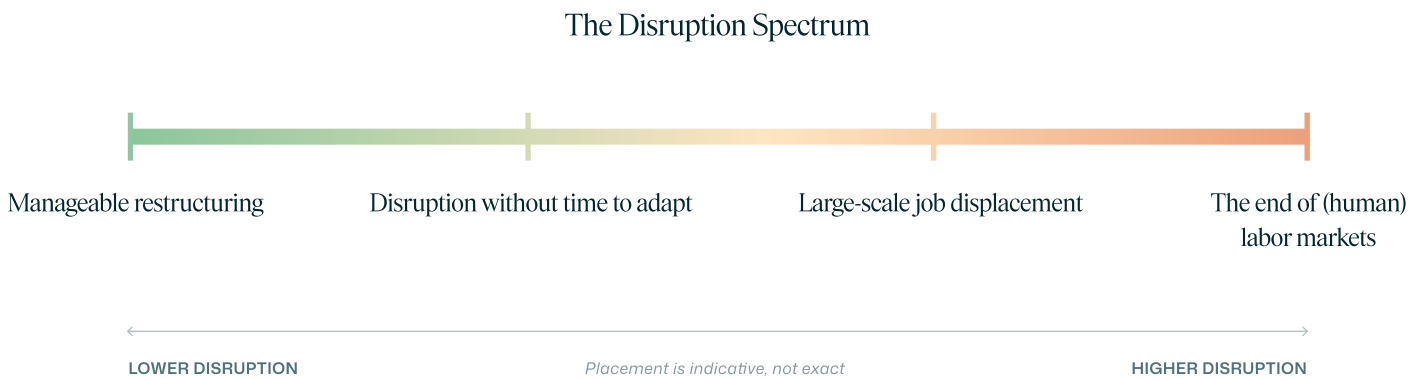
These gaps sit at the intersection of what governance most needs and what is hardest to measure with current infrastructure—meaning the data that matters most for action is the data we have least.

Who actually benefits from AI's gains?

A deeper gap underlies all of these: we don't have good ways to track who actually benefits from AI's gains. If those gains flow primarily to corporate profits and investors rather than wages, standard income statistics may miss the shift entirely—even as the share of growth reaching ordinary workers quietly shrinks. That matters because it's exactly the dynamic that would erode the tax base funding the safety net required to weather the transition. [Recent work by economists at the Peterson Institute](#) finds that national statistics were not designed to track AI's economic contribution, and warns that this measurement gap could become a policy gap if policymakers lack the data to decide who benefits from AI, who loses, and how the gains can be distributed.

What will actually happen to work?

This is an empirical question—or rather, a question that will eventually have an empirical answer. The honest position is that we don't know which of these futures we are heading into. The reality will likely be more complex than any single scenario: a mix of trajectories across sectors, regions, and worker groups. But four broad possibilities describe the range, from least to most disruptive:



1. Manageable restructuring

AI disrupts the mix of tasks and jobs—as electricity, computing, and the internet did before it—but the labor market absorbs the change over time. Some roles disappear; new ones emerge. Workers move, retrain, and adapt. Productivity gains eventually reach wages. The transition is uneven and sometimes painful, particularly for workers with fewer options, but no more severe than previous technology waves.

This is the scenario the historical record most directly supports. However, to be borne out, this scenario requires time, institutional support, and niches for human work: conditions that AI's distinctive features may undermine.

2. Disruption without time to adapt

AI displaces a meaningful share of knowledge and service work faster than new roles emerge to replace it. Productivity and profits flow disproportionately to capital and to a smaller, higher-skilled workforce. The job market splits: demand grows at the top, shrinks in the middle, and stagnates at the bottom, leaving fewer of the stable, mid-income roles that have historically anchored the middle class. Existing safety nets—designed for shorter spells of unemployment in a world where most workers can retrain into adjacent roles—prove inadequate for the scale and speed of the shift.

This scenario does not require AI to be uniquely transformative; it only requires it to be somewhat faster and broader than past automation waves, at a moment when institutions are already under strain.

3. Large-scale job displacement

AI proves capable of performing not just routine tasks but a wide range of cognitive work across many sectors. Labor's share of income falls significantly as AI capital substitutes for human labor at scale. GDP may continue to grow—even sharply—but the link between productivity growth and broadly shared prosperity breaks down. Work as the primary mechanism for distributing income becomes structurally inadequate. As employment shrinks, so does the tax base that funds the safety nets displaced workers depend on. This is the scenario where economic growth and worker wellbeing decouple.

4. The end of (human) labor markets

This scenario is the most speculative of the four: AI substitutes for human labor so comprehensively that employment ceases to be the primary way most people participate in the economy. Firms can produce at scale with minimal human input. Labor markets stop being the mechanism through which income, purpose, and social participation are organized. The challenge is no longer finding enough jobs or retraining displaced workers—it is renegotiating the social contract, and building entirely new systems for distributing resources and structuring lives.

What the scenarios share

Across all four, two things hold. First, the workers most at risk in the future are predictable today: those with fewer transferable skills, less geographic mobility, and the least access to institutional support. Second, any response—retraining infrastructure, income support, legitimate oversight bodies, and obviously the institutions that would underpin a new social order—take years, even decades, to build. Both are reasons to act before the scenario is clear. And the same is true of the opportunity: AI also creates genuine potential to redesign work for the better: more meaningful, more productive, more inclusive.

What should we do about it?

Scenarios describe what may happen. Narratives are the political and intellectual positions people hold about what AI's impact means, and what—if anything—government should do in response. The crucial distinction is that the same scenario can generate very different narratives. Understanding where these narratives sit matters, because they determine which policies get proposed, which actors support them, and where coalitions are possible or not. We group these positions into six broad narratives:

The Narrative Landscape



The Six Narratives

Ordered by assumed level of disruption, from least to most.

AI won't be that big. The disruption is overstated: firm-level AI adoption is slow, the occupational mix has shifted no faster than in earlier transitions, and worker anxiety is running ahead of actual labor displacement.

↳ *The counter-argument: Worker surveys already show widespread AI-related stress and fear—signals that may be leading, not lagging, indicators.*

Let the market sort it out. AI will disrupt work as past technologies did: some jobs will disappear, better ones will emerge. The main risk is overreacting: regulation and redistribution schemes slow the very innovation and job creation we're counting on to even out the disruption, and government workforce programs have a poor track record. The jobs that emerged from past technological transitions were largely unpredictable in advance, created by markets rather than policy design.

↳ *The counter-argument: It relies on conditions—time for markets to adjust, productivity gains reaching wages—that AI's speed and breadth may not allow.*

The Six Narratives

Ordered by assumed level of disruption, from least to most.

AI wealth will eventually reach everyone. AI will replace work, but also generate such extraordinary wealth that the question shifts from distributing jobs to distributing surplus. Material needs are met at low cost, therefore the central challenge becomes redesigning meaning and social organization, not poverty. This is speculative, but taken seriously by some of the people building the technology.

↳ *The counter-argument: Historical productivity gains haven't automatically reached workers. There's no mechanism that guarantees AI surplus flows downward without deliberate redistribution policy.*

AI will replace tasks—and we need private and public investment to help workers keep up. AI will make workers more productive, not obsolete—replacing tasks within jobs rather than jobs themselves, and creating new roles over time as past technologies did. But without active investment, that transition won't be fair or smooth. At the firm level, augmentation becomes intensification: closer monitoring, algorithmic management, more output without shared gains. At the system level, workers who need to retrain or move into new roles won't do so without serious public investment in reskilling, AI literacy, and sectoral training partnerships. The historical evidence supports this: retraining works when properly resourced and paired with income support.

↳ *The counter-argument: It assumes new jobs will emerge quickly enough—and in enough volume—to absorb the workers whose tasks disappear.*

Workers will be displaced by AI, so we need a safety net. In a significantly disrupted labor market, AI will substitute for a meaningful share of human labor, displacing workers, not just aspects of work. Not every displaced worker will find an equivalent job, therefore, we need policies that support those who won't. The response must go beyond retraining—through wage insurance, portable benefits, or ownership stakes in AI-deploying firms.

↳ *The counter-argument: If displacement proves slower or narrower than feared, expanded safety nets may reduce pressure to retrain and adapt without addressing the actual need: helping displaced workers move into new jobs.*

AI's gains must be distributed broadly. In the most transformative scenario, where employment itself stops being the primary route to livelihood, AI will leave behind even those who are employed. Without structural redesign—universal basic income, sovereign wealth funds, taxation of AI profits—gains will concentrate narrowly even as GDP grows.

↳ *The counter-argument: UBI and sovereign wealth funds are largely untested at scale and face enormous political and implementation barriers, making the case for them hard to build before the scenario is clear.*

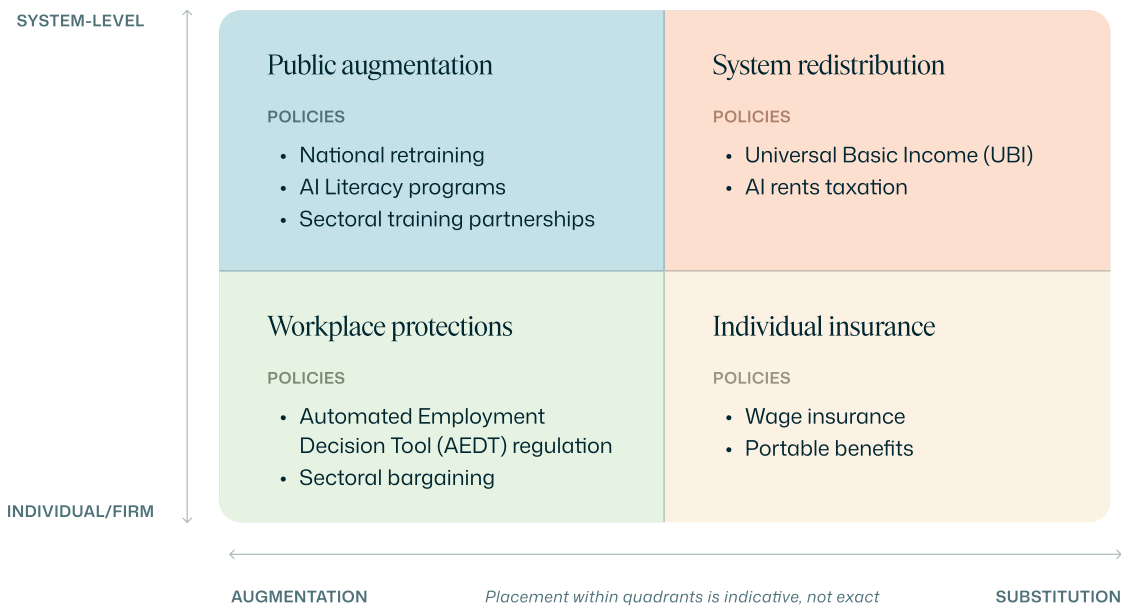
The six narratives are not mutually exclusive—most political or policy positions draw on more than one. But mapping where they sit helps clarify the different tools, actors, and timelines each calls for.

What acting looks like

Not everyone believes action is needed—and that narrative holds real influence. But for those who do, governance interventions can be organized along two axes:

- **Assumptions:** at one end, positions assuming AI will mostly augment workers (so policies framed around worker protection and reskilling will suffice); at the other, the belief that AI will substantially substitute for human labor (so the response must rebuild the social contract, not just protect employment and earnings).
- **Level:** whether intervention should happen at the individual or firm level (retraining, contract protections) or at the system level (federal programs, tax structure, redistribution).

The Policy Landscape



Augmentation × Individual/firm

Workplace protections: AI augments work, but workers need rights, voice, and training so augmentation doesn't become intensification or surveillance. Intervention happens at the workplace level, through tools such as automated-decision-tool regulation and sectoral bargaining. This is the most policy-active cell, and the area where the largest share of state AI legislation has been enacted.

Augmentation × System-level

Public augmentation: AI will augment workers, but only if public investment in retraining, AI literacy, sectoral training partnerships, and workforce institutions keeps pace with deployment. Without it, the benefits of augmentation flow disproportionately to employers. This cell is less politically active in the U.S. than the firm-level cell, but more developed in OECD, EU, and Nordic policy frameworks.

Substitution × Individual/firm

Individual insurance: AI will substitute for some labor, and the response is to help individual workers weather the transition through wage insurance, expanded benefits, and ownership stakes in AI-deploying firms. This cell is more theoretical than the augmentation × individual/firm cell but more politically accessible than system-level redistribution.

Substitution × System-level

System redistribution: AI substantially substitutes for human labor, and without system-wide redistribution (universal basic income, tax-base reform, sovereign wealth approaches, taxation of AI rents) gains concentrate at the top while wages and labor's share of income fall. This cell remains largely theoretical. Fathom's February 2026 polling found broad public support for workforce protections across the policy landscape: from familiar ones like retraining to more innovative proposals like a sovereign wealth fund, signaling that the Overton window is open.

Support for workforce policies

From Fathom's February 2026 report *What Americans Really Want*
Ranked by strongly favor

	STRONGLY FAVOR	SOMEWHAT FAVOR	TOTAL FAVOR
Retraining and re-skilling programs	50%	34%	84%
Transferable benefits that help workers easily switch between jobs	47%	36%	83%
Taxing companies that replace workers with AI to fund subsidized employment in "human-essential" sectors like health care, education, and infrastructure	45%	30%	75%
Universal basic income or universal basic services (such as free healthcare, childcare, and transit) so people need less income to live well	45%	29%	74%
Expanding existing protections like unemployment insurance and the earned income tax credit (EITC)	42%	35%	77%
Mandatory corporate policies that share gains with workers, such as employee stock ownership and profit-sharing requirements	40%	39%	79%
A sovereign wealth fund that creates public ownership of AI-related assets that pay dividends to citizens, so everyday people can benefit from AI's success	37%	34%	71%
New forms of worker organization, such as sector-wide unions or freelancer associations, that strengthen workers' collective bargaining power	35%	37%	72%

The Legislative Landscape

Enacted legislation is a lagging indicator—action plans and proposals tell us more about where the landscape is headed.

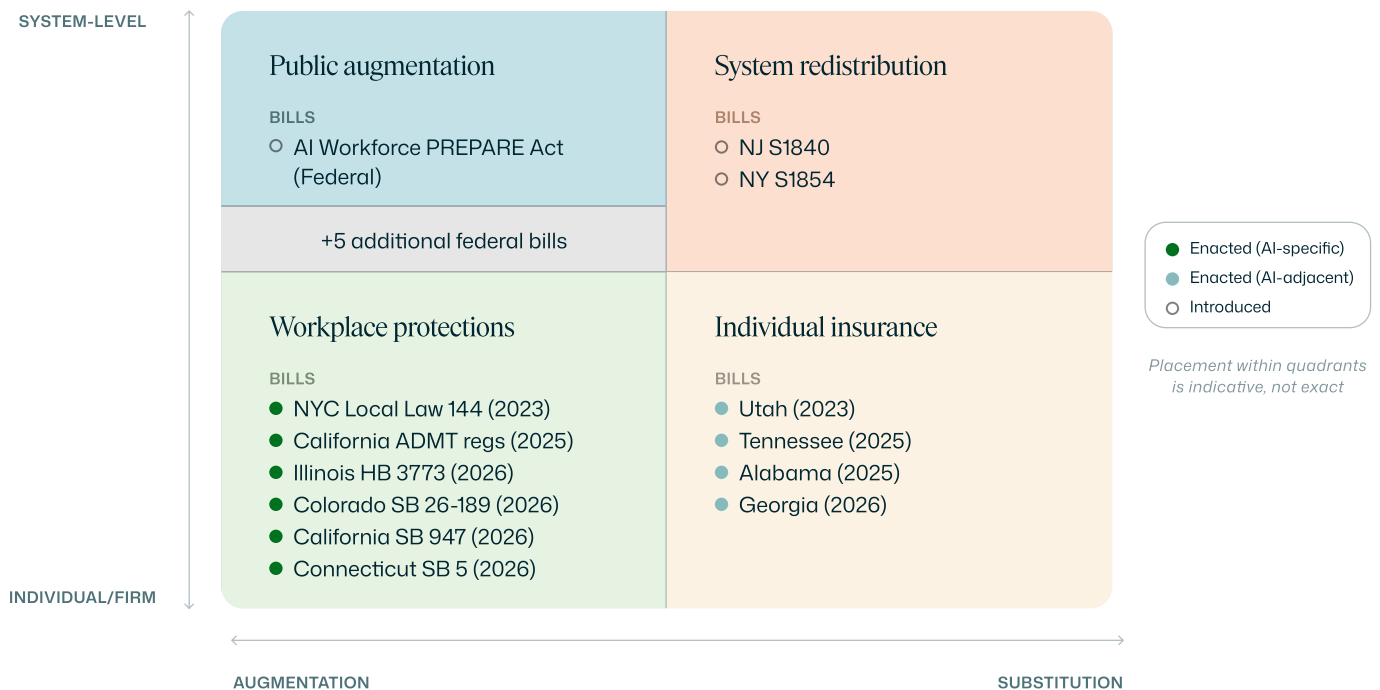
Six federal bills on workforce transparency and AI job-loss tracking have been introduced between March 2025 and March 2026, and the volume and ambition of proposals is likely to accelerate. The [AI Workforce PREPARE Act](#) (introduced December 2025) is among the most significant bipartisan federal signals: it would require employers to flag AI-driven layoffs and direct the Department of Labor to produce AI-specific workforce forecasts. Others focus on worker training grants, employer tax credits for AI literacy programs, and mandatory employer disclosure of AI-driven hiring and layoff decisions—reflecting a federal landscape that is moving, but still concentrated on measurement and workforce preparation rather than structural adjustment measures.

At the state level, as of early 2026, more than 1,500 AI bills had been introduced across 45 states; of the roughly 1,200 from 2025, 145 became law. Employment is the only category where a significant share has passed—but almost entirely on one side of the policy landscape. A few state bills, including in New Jersey ([S1840](#)) and New York ([S1854](#)) point further toward the system level, proposing to channel revenue from AI-deploying firms toward worker support.

Overall, enacted legislation clusters in measurement and workplace protections: algorithmic accountability, bias audits, hiring transparency, and surveillance limits. Legislation on the substitution side (wage insurance, retraining mandates, redistribution) is limited as of yet.

The Legislative Landscape

as of June 2026



How the rest of the world is responding

Outside the U.S., workforce policy in response to AI is emerging along three broad lines: regulatory emphasis and social welfare in Europe, growth-plus-stability balancing acts in China, high-adoption Asian economies, and parts of the wider Global Majority.

European Union and Nordics: regulation plus social protection. The EU AI Act treats all AI systems used in employment and hiring as “high-risk,” imposing obligations on providers and users including risk management, human oversight, and bans on applications such as emotion recognition in workplaces. These high-risk obligations, originally set for August 2026, are now deferred to December 2, 2027 under the Digital Omnibus agreement reached in May 2026. In Nordic countries, AI-specific labor provisions sit on top of strong collective bargaining, generous unemployment insurance, and active labor-market policies—the institutional environment that proponents of distributional mechanisms often have in mind when they argue the transition is manageable.

China and high-adoption Asian economies: rapid automation with stability concerns. China is automating aggressively to offset demographic decline and boost productivity, having installed roughly half the world’s industrial robots, and manages the resulting displacement risk through general labor law, pressure on firms to avoid mass layoffs, and case-by-case arbitration—but without a dedicated AI-and-work framework.

Global Majority: high exposure, limited buffers, early experiments. Many countries in the region face high automation risk with fewer institutional safety nets: the ILO estimates more than half of jobs in low- and middle-income countries are at high risk. Responses so far are small-scale and experimental (digital- and AI-literacy programmes and basic-income pilots) rather than parts of an integrated AI transition strategy.

What's still missing

The gaps in the future of work landscape are not a matter of lack of ideas—the policy menu is rich and legislation is moving. The harder problem is more fundamental: there is not yet agreement that action is needed, a clear signal of what would trigger it, or institutions trusted to lead it. These are the gaps that matter most:

The consensus gap

There is no agreement on whether to act. The deeper problem is political, not technical: the narratives in this report reflect genuine disagreement—not just about what policies to pursue, but about whether government should act at all. That disagreement is unlikely to resolve before decisions need to be made. Building consensus—through better evidence and forums that aim to create consensus among actors—is itself a governance priority, and a very foundational one.

The translation gap

Research isn't becoming policy. Research institutions are getting better at mapping exposure and modelling scenarios, but few concrete mechanisms translate those findings into triggers, guardrails, or obligations for firms. Most proposals remain at the level of frameworks, menus, or principles rather than deployable institutional arrangements.

The data gap

Measurement: National statistics weren't designed to track AI's economic contribution—there are no agreed leading indicators for AI-related layoffs, entry-level hiring shifts, or wage compression, and distributional impacts are rarely broken down with the geographic and demographic granularity that would move policymakers to act. The data that matters most for early action is the data we have least.

Data access: evidence isn't shared and can't be independently checked. Technology companies and large employers increasingly hold the most detailed data on how AI is used at work, while independent researchers, policymakers, and civil society often see only lagging or fragmented statistics. The UK's newly established AI Economics Institute, whose function is to build the evidence base on AI's economic and labor market impacts, is an early attempt to address exactly this—with data-sharing agreements already in place with LinkedIn and collaborations agreed with Anthropic, Google, OpenAI, and Microsoft. Wider, more standardized access would serve two ends: helping independent bodies validate claims about augmentation, job creation, or displacement, and giving policymakers the evidence to design responses around what is actually happening.

The prototyping gap

Ideas aren't being tested. The field has produced a rich set of ideas for AI and work—data trusts, worker-centric impact assessments, new social-insurance models, new models for continuous and modular learning, job-quality guarantees—but few have moved from reports into practice. Piloting and iterating them is the natural next step, and would reveal which institutional forms actually hold up.

The trust gap

Narratives aren't anchored in evidence. Polling shows that people are already worried and conflicted about AI and work, and public narratives often run ahead of what labor-market data actually shows. Better research connecting public sentiment to that evidence would help ground the conversation in public mandate.

Foundations for any future

The following cross-cutting solutions are largely scenario-agnostic. They do not depend on resolving the augmentation-versus-substitution question first. That is precisely what makes them urgent: they are the foundations every future will need, and they take years to build.

- ↳ **Anticipatory indicators and governance triggers.** There's little agreement on what would signal that AI is disrupting the labor market faster than institutions can keep up, leaving policymakers to rely on lagging quarterly statistics. A valuable step would be to bring policymakers, researchers, labor, and industry together to agree on a small set of leading indicators—such as AI-related layoffs, shifts in entry-level hiring, or wage compression in exposed occupations—that could flag the need for attention before damage accumulates.
- ↳ **Independent measurement and data-sharing.** Because the most decision-relevant data sits inside private systems, two needs stand out: trusted arrangements for sharing AI usage and labor-impact data across firms, labs, researchers, and governments, and independent mechanisms to validate claims about augmentation, displacement, and reskilling—neither of which requires exposing sensitive information. And since the party deploying a technology cannot also be the one validating its labor effects, the gap is not measurement alone but verification, where models already used elsewhere in AI governance may be adaptable.

SPOTLIGHT

Open Banking and the utility question: Can we mandate data-sharing?

The challenge of compelling private actors to share data with independent parties has precedent. In 2016, the UK's Competition and Markets Authority mandated that the country's largest banks share customer data with third-party fintechs—a model that has since spread internationally. The U.S. attempted something similar through the CFPB's open banking rule, but it has stalled under legal challenge, with banks arguing the requirements impose undue burdens.

The analogy is imperfect: financial data and AI deployment data are different, and the purposes would be distinct: the UK mandate was aimed at innovation, whereas mandatory AI data-sharing would serve policymaking. But the underlying governance question is the same—can regulators compel private actors to share data that serves a public interest? Open banking shows it is possible. It also shows it is hard, even in financial services, which is among the most heavily regulated sectors in the economy. That the U.S. couldn't get it done in banking is a cautionary note— but also suggests that if the political will exists, the policy architecture is buildable.

One emerging frame in this space is whether frontier AI systems could eventually be classified as public utilities—which would carry much stronger data-sharing and access obligations. At BlackRock's Infrastructure Summit in March 2026, Sam Altman suggested that [artificial] intelligence is a utility like electricity or water and people buy it from us on a meter—he meant it as a business model claim, but one that, if taken seriously, would have significant regulatory implications. This remains a relatively fringe conversation, but one that could gain momentum in the coming years—especially as states are beginning to regulate data centres the way they regulate utilities.

What We Can Start Building Now

- ↳ **Piloting institutional forms.** Many promising ideas—data trusts, verification mechanisms, new social-insurance and learning models—remain on paper. Moving a few from proposal to small-scale pilot, and iterating on what works, would let the field learn which forms hold up before they are scaled—and is where independent, non-commercial actors are especially well placed to lead.
- ↳ **Bridging institutions and convening.** Connecting macro debates about substitution and redistribution to the more immediate questions of labor-market governance, public legitimacy, and political feasibility will require convening actors who currently operate in separate spaces, including bridging the divide between workforce policy and the broader AI-governance community.
- ↳ **Preparedness and education.** Beyond retraining for today's jobs, a longer-horizon priority is preparing people for an economy where the mix of valued skills keeps shifting. That points toward a change in education itself—away from a front-loaded model, where formal qualifications are completed once early in life, toward something more continuous: modular, recurring, and available across a whole career. Building that adaptability is low-regret whether AI proves transformative or merely significant.
- ↳ **Better public evidence on trust and legitimacy.** Existing polling consistently shows that Americans are worried about AI's impact on work and open to intervention, but it rarely answers the most governance-relevant question: who do people trust to act? Mapping support for specific interventions against trust in employers, government, and independent bodies would reveal where public support is high but no trusted actor exists to act on it, critical input for both policy design and implementation.

AI's workforce disruption could also be an opportunity to redesign work itself. Automating tasks that are dangerous, repetitive, or low-value could redirect human effort toward higher-value contribution—and governance proposals should be open to that possibility, not just focus on protecting the status quo.

Conclusion

The future of work is one of the fastest-rising issues in the AI debate—and one of the most consequential. It is not only an economic question. Work is how most people structure their days, develop competence, connect with others, and answer the question of who they are. What AI does to work will shape not just livelihoods but the texture of human life.

There is no agreement on what will happen. The measurement is fragmented, the scenarios are wide-ranging, and the narratives that shape policy responses reflect genuine disagreement—not just about what to do, but about whether to act at all. That uncertainty is real and should be taken seriously. But it is not, alone, a reason to wait. The window for building the right governance infrastructure will not stay open indefinitely, and the risks of inaction are asymmetric: the workers most likely to be left behind are already predictable, and the institutions needed to support them take years to build.

The path forward does not depend on settling the augmentation-versus-substitution question first. The most valuable work that we suggest—better public evidence, independent measurement, anticipatory indicators, shared data infrastructure, bridging institutions—is robust across scenarios and worth building whether disruption proves large or small. We expect 2027 and 2028 to begin revealing which future we are in. The window for preparing is now, not then. The goal is not just to protect workers from what AI might disrupt, but also to ensure that what AI could make possible is broadly shared.

Fathom plans to map this landscape and to work with others toward a shared framework for action. The scenarios and narratives here are tools for dialogue and an invitation to start a critical national conversation. Building the right governance infrastructure now, while the political and technological picture is still taking shape, is the surest way to ensure that whatever workforce scenario plays out, it is one society can see clearly—and shape deliberately.



ABOUT US

We find, build, and scale practical solutions designed for the AI century. Fathom is an independent nonprofit whose mission is to build a global architecture that helps society navigate the transition to a world with AI by fostering trust, safety, and innovation.

Learn more at: fathom.org
Follow us: linktree/fathom.org

Interested in diving deeper into our research and insights?

Reach out to us at info@fathom.org to learn more.