ANTHROP\C

Anthropic's Economic Policy Framework

June 2026

We're still early in AI's adoption, but evidence from the US already suggests that entry-level workers in occupations most exposed to AI have seen weaker employment growth in recent years.[1] Economists' models of the future range from modest productivity-driven growth[2] to scenarios where economic output doubles but wages collapse.[3]

If AI capabilities continue to advance, AI may become a general substitute for human labor, with effects on work that are broad and, for many, difficult to absorb. At the same time, AI could generate economic growth on a scale the modern world hasn't seen. The central challenge isn't how to stimulate growth; it is making sure the gains are widely shared.

1 Brynjolfsson et al (2025), "Canaries in the Coal Mine"; Massenkoff and McCrory (2026), "Labor market impacts of AI: A new measure and early evidence"

2 Acemoglu (2024), "The Simple Macroeconomics of AI"

3 Korinek and Suh (2024), "Scenarios for the Transition to AGI"

Our values

We are not seeking job displacement. We are working to prevent or minimize it. Some amount of displacement, though we cannot say how much, may be an intrinsic consequence of the technology, and our responsibility is to prepare for it and respond to it. That is what this framework is for.

Supporting people financially is necessary, and it is the primary focus of this document because it is the more tractable of the two problems. It is not sufficient. There is dignity in work. We should help people find work wherever we can, and society should keep searching for ways to remain at full employment.

At the same time, people are more than their jobs. Over a longer horizon, AI may move society toward a world in which people work far less, or in which work carries a different meaning than it does today. Changes of that scale would reach beyond labor policy into how society is organized and how it relates to work.

Whatever happens, we are on the side of people. We are trying to solve these problems. We take no satisfaction in contributing to them, and we are not working to make them more likely.

About this framework

This document sets out an initial framework for US policy responses, organized by how severe the impact on jobs turns out to be. It separates the steps that make sense now in almost any scenario from the responses worth having ready if conditions worsen, and from the harder questions where research and dialogue should inform decisions.

We can't forecast with confidence where AI capabilities will be even a few years from now. The recommendations here are calibrated to what we expect from the broad diffusion of AI capabilities that exist today or will arrive within roughly the next year.

In the face of this uncertainty, our aim is to buy time—month by month, and year by year—by making the economy more flexible and resilient as AI accelerates, so that workers, firms, and governments have room to adapt. Whether an AI-fueled economy offers people new kinds of work, and whether workers can actually reach those opportunities, depends on the choices we as a society make. The economy has real capacity to adapt, but adaptation is not automatic.

This document is focused on public policy, but governments are not the only actors whose choices matter. Firms, including Anthropic, will shape this transition by how they deploy AI, and where possible they should choose the path that gives workers and the economy room to adapt. AI can be used to produce the same output with fewer people, or to accomplish more with the same number of people: building new products, serving markets that were previously out of reach, and raising the capability of less-experienced workers.

There are concrete steps firms can take. Companies adopting AI can build workforce training into deployment; decide in advance how freed-up capacity will be used, rather than defaulting to headcount reductions; retrain and redeploy people as roles change; and redesign early-career roles around AI. We are trying to make that path easier to take by providing training and enablement to our business customers alongside the technology, encouraging deployments that build worker capacity, giving them tools to measure their

workforce's growing fluency, and building products that make the people using them more capable, not just faster.

Yet adaptability has limits. If AI becomes a general substitute for human labor across most of the economy, realistic combinations of retraining, job matching, and mobility policy may fail to deliver a stable, full-employment equilibrium.

In this case, policymakers will need to carefully consider the pace of AI rollout. When technological progress causes large disruption and the available redistributive tools are imperfect, economic theory suggests the technology's rollout should be steered, or sequenced.[4] In normal times, the economy is built to absorb disruptions. But if disruption happens beyond the economy's ability to adapt, at the levels described in Tier 2, we would support government regulations and incentives at the firm level that manage the pace of displacement. Uniform application across the field is essential: any firm that slows down alone simply cedes the market to those that do not.

Each tier buys time for workers and the economy to adapt, and eventually judge whether the interventions in the next tier are needed. More difficult, longer-term questions around what happens if the link between work and income breaks are addressed in the highest-disruption tier.

While this particular proposal is focused on the US, the principles underpinning this framework are intended to be global. Preparing institutions ahead of disruption, sharing AI's gains broadly, and modernizing the systems workers depend on will be necessary anywhere AI touches. We recognize that the right policy mix in any country will turn on its labor law, social contract, and where models are developed. We have started with the US because Anthropic is headquartered in San Francisco, and because Claude is used more in the US than any other country.[5]

4 Korinek and Stiglitz (2026), "Steering Technological Progress"

5 Appel et al (2026), "Anthropic Economic Index report: economic primitives"

The foundation

The framework that follows depends on three foundations: the ability to see AI-driven labor market disruption as it develops, the analytical capacity to interpret what we see, and the delivery infrastructure to act on it.

Measurement. Governments' statistical infrastructure was built for slower-moving economies; it does not yet track AI adoption or its labor-market effects with the speed or granularity this moment requires. We need expanded investment in our statistical agencies, sustained funding for AI-specific survey instruments, and reporting requirements for AI labs and firms—including Anthropic—on deployment patterns and workforce effects.[6] Instruments should focus on directly observable metrics like AI usage rates and intensity of use, and worker productivity indicators by level of AI adoption.

Anthropic is contributing to this effort through the Anthropic Economic Index, which uses Claude usage data to track how AI is being used across occupations and industries.[7] We believe transparency about AI's economic effects is a precondition for informed policymaking, but data from a single company cannot tell the whole story.

A government unit focused on tracking AI's effects. Visibility is necessary but not sufficient. Governments also need a small, dedicated unit to track how AI is moving through the economy sector by sector, stress-test the institutions that will have to absorb the shock, and flag the early signals that should trigger a policy response. Governments have built this kind of cross-cutting analytical function before—the Council of Economic Advisers was created for the post-war transition—and we need an equivalent for AI.

Delivery infrastructure. The delivery infrastructure for Unemployment Insurance (UI), which provides workers time to navigate job loss and transition,

6 The American Economic Association Committee on Economic Statistics (2025), "Measuring the Economic Effects of Artificial Intelligence"

7 Anthropic (2026), "Anthropic Economic Index"

is not sufficiently prepared for a large labor market shock.[8] Many states still run UI on legacy systems that buckled under pandemic claim volumes, took months to implement emergency extensions, and made billions of dollars worth of improper payments.[9] Modernized systems could process claims faster, flex quickly when Congress or states adjust benefits, and run real-time identity and cross-state fraud matching.

Only 52% of the world's population is covered by any social protection scheme,[10] and in many countries the delivery systems that do exist cannot absorb a fast-moving labor market shock. Countries should invest now in the rails required to expand and adjust coverage quickly. The same logic extends to revenue collection and the underlying digital public infrastructure—identity, payments, data exchange—that any of these programs would run on.

The policy framework

We cannot know today which scenario will unfold, but we believe this uncertainty warrants preparation. We need to invest now in measures that make sense in any scenario, and design in advance the responses we would want ready if disruption deepens. For this reason, our policy recommendations anchor on three possible scenarios.

We suggest calibrating responses to the severity of AI-driven labor market disruption, using the unemployment rate as the primary trigger because it is widely understood, regularly measured, and directly captures labor market health. But the unemployment rate is not the only thing that matters. People may keep their jobs while the pay, security, and quality of that work erode. Governments should therefore watch a broader set of signals alongside unemployment: labor force participation, underemployment, wages, and labor's share of national income.

8 Ganong et al (2022), "Lessons Learned from Expanded Unemployment Insurance during COVID-19"

9 US Government Accountability Office (2022), "Unemployment Insurance: Transformation Needed to Address Program Design, Infrastructure, and Integrity Risks"

10 International Labour Organization (2024), "World Social Protection Report 2024-26"

We also cannot know for certain what form disruption will take. One possibility is a temporary shock, after which the labor market re-stabilizes. The other is an enduring restructuring, in which the demand for human labor is significantly and persistently lower. These are different problems. A temporary shock calls for temporary supports withdrawn as the labor market recovers; an enduring restructuring calls for more permanent responses.

Our suggested interventions serve one or both of two aims: stability and adaptability. Some help keep the macroeconomy stable, through income support and transfers that cushion the shock and sustain demand. Others build the economy's adaptability, by moving people toward more durable work, easing transitions between jobs, and, where possible, slowing the pace of displacement. Neither is enough on its own, and our key interventions do both: unemployment insurance, for example, stabilizes demand while giving displaced workers the time and tools to find a well-matched job. As disruption increases, the balance shifts from adaptability toward income support and redistribution.

The policies within each tier are those we currently find most promising on our reading of the existing evidence; we expect our thinking to evolve as more accumulates. Some measures that make sense in a near-term, transitory shock could make matters worse in an enduring restructuring by entrenching the link between income and employment that a Tier 3 world may need to loosen. We flag those measures below, and we expect to phase out some Tier 1 and Tier 2 tools if conditions move toward Tier 3.

Tier 1: Baseline economy (~5% unemployment with churn)

Prepare for disruption and ensure families share in AI's upside

At 5% unemployment, the headline numbers look normal. But underneath, the labor market may be churning faster than workers can adapt. AI is automating tasks across occupations simultaneously, steadily compressing the value of skills that workers spent years acquiring.

Establish universal pre-distributive capital accounts. We should give every American a direct financial stake in the economy at the moment when AI is fueling growth. The federal government has already taken a step in this direction with new capital accounts seeded at birth. We believe these accounts should be made permanent and expanded towards universal coverage, with initial priority given to the cohorts most exposed to near-term AI disruption:

- Phased eligibility expansion: beyond children, accounts should be extended to young adults entering the workforce and to incumbent workers in occupations most exposed to AI-driven displacement, with a defined path to coverage for all working-age adults.
- Flexible withdrawal for workforce transitions: account holders should be able to draw on funds for retraining, relocation, credentialing, and other transition expenses.
- Equity funding: policymakers should expand the mechanisms by which these accounts can be funded, including with equity in AI companies, so that beneficiaries share directly in nearer-term gains from AI-driven growth.

To be clear, this measure does not cushion near-term displacement. We propose it in Tier 1 because, for the intervention to matter, it must start before disruption is visible. Accounts compound; the earlier they are seeded, the more they are worth when they are needed. We recognize that pre-distributive accounts are slow-compounding and cannot alone compensate for permanent job loss that happens soon—they must be paired with the income support and transition mechanisms in Tiers 2 and 3.

Additional Tier 1 interventions

INTERVENTION	DESCRIPTION
Wage insurance	Time-limited, capped supplemental income for workers who take lower-paying jobs after displacement, funded by the government. US evidence suggests wage insurance decreases nonemployment and increases employment rates such that it is funded by avoided UI payments and increased income tax.[11] Evidence from studies in Canada and Germany is less encouraging.[12]
Occupational licensing reform	Federal funding for states that make it easier to enter licensed occupations, for example by recognizing licenses issued in other states or shortening the path to qualification, without weakening the rules on what licensed workers may do or who supervises them. Roughly 24% of US workers hold a certification or license[13] and evidence shows licensing requirements reduce employment and raise prices without corresponding net welfare gains.[14]
Retention tax incentives	Multi-year tax credits for companies that retain and redeploy workers into substantially different roles when adopting AI. Such incentives can reduce distributional harm[15] and preserve employment during shocks, but can also prop up low-productivity firms.[16] Design would need to overcome implementation challenges including noncompliance[17] and fraud.[18]

11 Hyman et al (2024), "Wage Insurance for Displaced Workers"

12 Bloom et al (2001), "Testing a Financial Incentive to Promote Re-Employment among Displaced Workers: The Canadian Earnings Supplement Project (ESP)"; Brussig et al (2006), "Entgeltsicherung — Befunde einer Evaluation"

13 Bureau of Labor Statistics (2026), "Certification and licensing status of the employed by occupation, 2025 annual averages"

14 Kleiner and Soltas (2023), "A Welfare Analysis of Occupational Licensing in US States"

15 Lehr and Restrepo (2023), "Optimal Gradualism"; Beraja and Zorzi (2025), "Inefficient Automation"

16 Giupponi and Landais (2023), "Subsidizing Labour Hoarding in Recessions: The Employment and Welfare Effects of Short-time Work"

17 US Government Accountability Office (2003), "The Worker Adjustment and Retraining Notification Act: Revising the Act and Educational Materials Could Clarify Employer Responsibilities and Employee Rights"

18 Beebe (2023), "A Lesser-Known Pandemic Relief Program Inspires Widespread Fraud"

Additional Tier 1 interventions (continued)

INTERVENTION	DESCRIPTION
Workforce training grants	Grants for employer-connected sectoral training programs that place workers into roles with strong current demand. Evidence is mixed on traditional government retraining,[19] but shorter-term programs connected to specific employers and sectors have shown significant income gains for participants.[20]
Job matching infrastructure	Invest in AI-powered skills-based labor market platforms, Learning and Employment Records, and skills wallets. Rigorous evaluations of job recommendation systems in a number of countries find modest effects on reemployment.[21]

Tier 2: Recession-level disruption (~10% unemployment)

Expand temporary income support and retraining

At 10% unemployment, disruption is a fact of life for tens of millions of American families. Professions no longer need the expertise they once demanded, careers end earlier than planned and begin later than expected, job searches stretch from months into years. Families revise plans—for homes, for children's education, for retirements, or abandon them. Communities that depended on a stable base of steady employment find that base eroding faster than new sources of work can replace it. Familiar pathways—from education to career, from effort to advancement, from work to stability—have become less reliable, and for many, impassable. The question shifts from whether to intervene to how fast existing systems can scale.

This framework's aim is to use the safety net as a launchpad to help workers pursue new vocations. This will require investing, in parallel, in the opportunities and industries where meaningful work will be.

19 Bloom et al (1997), "The Benefits and Costs of JTPA Title II-A Programs: Key Findings from the National Job Training Partnership Act Study"; Card et al (2018), "What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations"

20 Katz et al (2020), "Why do sectoral employment programs work? Lessons from WorkAdvance"

21 Belot et al (2019), "Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice"; Behaghel et al (2025), "The Potential of Recommender Systems for Directing Job Search: A Large-Scale Experiment"; Le Barbanchon et al (2023), "How can AI improve search and matching? Evidence from 59 million personalized job recommendations"

Expand unemployment insurance. UI is the most direct, tested, and rapidly scalable mechanism for supporting displaced workers. The Extended Benefits program established in 1970 extends UI when state unemployment is high, but its triggers can fire late and unevenly, and Congress has repeatedly had to legislate emergency extensions on top of it. We should consider reforming the trigger, so that extension is automatic and uniform across states.[22] Macroeconomic evidence supports extending the duration of UI during periods of elevated unemployment: during the Great Recession and COVID pandemic, extended benefits did little to discourage people from looking for work, while helping families keep up with essential spending.[23, 24]

Additional Tier 2 interventions

INTERVENTION	DESCRIPTION
Sector-specific transition support	Targeted assistance for workers in industries facing acute AI-driven displacement. Could include extended income support, subsidized retraining, job search allowances, and relocation assistance. Evidence suggests that intensive reemployment services, including career counseling and job search assistance, produce modest positive effects on earnings.[25]
Basic needs relief	Monthly payments to displaced workers who have exhausted UI benefits or labor market entrants unable to secure employment after some period of documented search, ensuring continued access to food, housing, and healthcare. Enhanced benefit levels for those who opt into roles addressing labor shortages in healthcare, education, child and elder care, infrastructure, ecosystem restoration, and public safety. Research shows consumption drops sharply when workers exhaust UI benefits,[26] and existing safety net programs only partially fill the gap.[27] This relief is structured to reach a broader population and could be delivered either as a new instrument or as an expansion of those existing programs.

22 Chodorow-Reich and Coglianese (2019), "Unemployment Insurance and Macroeconomic Stabilization"

23 Chodorow-Reich et al (2019), "The Macro Effects of Unemployment Benefit Extensions: a Measurement Error Approach"

24 Ganong et al (2021), "Micro and Macro Disincentive Effects of Expanded Unemployment Benefits"

25 Mathematica (2008-17), "Workforce Investment Act Adult and Dislocated Worker Programs Gold Standard Evaluation"

26 Ganong and Noel (2019), "Consumer Spending during Unemployment: Positive and Normative Implications"

27 Rothstein and Valletta (2017), "Scraping By: Income and Program Participation After the Loss of Extended Unemployment Benefits"

Tier 3: Transformative disruption (unemployment exceeding historical peaks)

Broaden and deepen existing support while building new mechanisms for redistribution

At this level of disruption, we are past the edge of the maps that policymakers and economists have historically used to navigate: unemployment at levels never before sustained alongside an economy generating record output. The search for work stretches past a year, then past two, and for some, eventually stops. Savings built over a working life are drawn down; rent is paid late, then not at all. The link Americans have long taken for granted—between contributing to the economy and sharing in its rewards—is strained or broken.

The policy challenge shifts from temporary transition support to sustained income replacement for a large share of the workforce. We are less certain about the right answers here than in earlier tiers, and we expect to develop this thinking in partnership with The Anthropic Institute, policymakers, and external researchers. What follows is our current assessment of the directions most worth exploring.

Sustain income through existing systems. We should consider automatically extending UI duration as labor market conditions deteriorate. That extension should come alongside enhanced wage insurance with higher caps and longer duration, and expanded basic needs relief with broader eligibility, covering both underemployed workers and those who never qualified for traditional UI. As extensions lengthen, benefit levels should converge toward a common floor rather than continuing to track prior wages indefinitely. When displacement stretches into years, paying people in proportion to what they used to earn is harder to rationalize on fairness or fiscal grounds, and a common floor creates a natural on-ramp to the mechanisms below.

In the case of a rapid and large rise in unemployment, deficit-financed income support is textbook countercyclical policy. But if AI delivers transformative productivity gains, the aggregate economy may grow substantially alongside sustained technological displacement—an unemployment scenario that is more structural than cyclical. This calls for policies that are likewise structural and fiscally sustainable. Historically, rapid productivity growth has expanded the tax base and improved fiscal positions. However, the US may enter this period with

elevated deficits and recent inflationary experience, limiting fiscal headroom for large-scale interventions. And if displacement instead coincides with continued inflation or fiscal crisis, the toolkit would need to differ—prioritizing reforms that don't require new spending, such as regulatory adjustments, private insurance mandates, or reallocation within existing budgets.

Establish new tax bases and redistribution mechanisms. If AI capabilities advance rapidly and broadly, we may face labor market conditions without historical precedent—not merely high unemployment, but a structural decoupling of productive contribution from income. Because effective tax rates on labor substantially exceed those on capital, a shift in national income from labor toward capital means the current tax system captures a shrinking share of a growing economy, even as the cost of sustained income support rises.[28] We will need new revenue sources. Traditional safety net tools may not be sufficient.

We are not yet ready to advocate specific policies for this scenario. But we can name a few candidate mechanisms that may become necessary and invest in the research required to evaluate them:

- Potential revenue sources could include increasing the capital gains tax, broad-based consumption taxes, sector-specific levies on AI use (measured by tokens, compute, or revenue), and scalable "digital dividends" funded by taxes on the digital sector.[29]
- Potential redistribution mechanisms could include universal basic income, AI sovereign wealth funds funded by investment stakes in AI-driven productivity, equity-sharing mechanisms giving workers partial ownership in AI enterprises, and dramatically expanded pre-distributive capital accounts building on existing models.

Substantially expand public investment in human- and community-facing work. Some of the roles that may prove more resistant to AI substitution are also where the US faces well-documented, persistent workforce shortages.[30] Directing a portion of any AI-driven surplus toward this work serves two

28 Korinek and Lockwood (2025), "Preserving Fiscal Stability in the Age of Transformative AI"

29 Marinescu (2025), "Resilient by Design"

30 US Government Accountability Office (2022), "K-12 Education: Education Should Assess Its Efforts to Address Teacher Shortages"; US Health Resources & Services Administration (2025), "Health Workforce Projections, 2023–2038"

goals: it closes staffing gaps that current funding levels have not, and it keeps paid work available for those who want it. We recognize the second rationale rests on a value judgment—that work confers something income alone does not—which not everyone shares and which future generations may weigh differently. The case here does not depend on it.

We're actively studying these mechanisms, and we plan to invest in the research required to better evaluate them. Some are more promising than others, and some carry distinct trade-offs. For example, broad-based approaches offer simplicity but may lack precision, while displacement-indexed mechanisms are more targeted but harder to administer.

Regardless of the tax base or distribution mechanism: we are ready and willing to pay our fair share. We believe that if AI companies generate transformative returns, they have an obligation to ensure those returns are broadly shared.

We have significant agency to shape how the future unfolds, even with significant uncertainty about the curve ahead. If we do face an economy in which most income no longer comes from work, we expect society's responses will necessarily rely on help we do not yet have, including from increasingly capable AI systems themselves. For now, our commitment is narrower and concrete: to be specific about what we propose to do at each step along the path, to fund the research that tells us whether it is working, and to revise as the evidence and the technology change.