U.S. AI LABOR-MARKET TAIL RISK

Stress test for persistent displacement, underemployment, wage compression, and safety-net demand, 2026-2040

Prepared for a domestic policy agency
Date: 16 March 2026
Purpose: size labor-market buffers for worst plausible AI-disruption outcomes rather than extend ordinary business-cycle precedent.
Use: policy planning stress test, not a base-case forecast.

2035 severe tail	2035 extreme tail	2035 wage pressure
Support population 17.6m	Support population 31.8m	Low-barrier gap -21% extreme
Official unemployment 7.4%	Official unemployment 10.1%	Mid-barrier gap -15% extreme
LFPR 58.3%	LFPR 55.2%	High-barrier gap -6% extreme

Bottom line. The main planning risk is not unemployment alone. In both tail scenarios, AI drives a combined stock of unemployed, underemployed, and detached workers that is materially larger than the official unemployed count. The 2030-2038 window is the most acute because capability growth, enterprise deployment, and worker reallocation friction all reinforce each other.

This document intentionally leans forward: it treats current BLS labor-market structure as the starting point, but it does not assume that future AI disruption will resemble past recessions or past sectoral transitions.

Executive summary

1. Historical-cycle models likely understate tail risk. AI capability is improving much faster than normal labor-market retraining cycles, enterprise deployment is already broad and increasingly automation-oriented, and early hiring effects are already visible in exposed entry-level work. Current aggregate unemployment data therefore give a false sense of safety.

2. The policy-relevant support population can become very large. By 2035, the model yields about 17.6 million people needing support in the severe planning case and 31.8 million in the extreme plausible case. Those totals combine AI-linked unemployed, underemployed, and detached workers.

3. Wage pressure in the remaining human-employable task bands is likely to be severe. By 2035, modeled wage level gaps versus the no-tail-shock reference reach roughly -11% / -8% / -4% in the severe case and -21% / -15% / -6% in the extreme case across low-, mid-, and high-barrier human tasks.

2035 planning metrics

Metric	Severe tail	Extreme plausible
Official unemployment rate	7.4%	10.1%
Modeled underutilization rate	12.8%	20.2%
Labor-force participation rate	58.3%	55.2%
Ever displaced workers	33.8m	48.6m
AI-linked unemployed	5.2m	8.8m
Long-term AI unemployed	2.9m	6.0m
AI-linked underemployed	5.5m	8.6m
AI-linked detached / discouraged	6.8m	14.4m
Total support population	17.6m	31.8m

Interpretation note. Official unemployment is only part of the burden. The modeled underutilization rate includes AI-linked underemployment and a partial weight on detachment to reflect discouraged workers and unrealized labor supply. Wage gaps are cumulative versus the reference path, not nominal wage cuts.

Why backward-looking models likely understate AI tail risk

■ **Capability progress is fast enough to create step-change risk.** METR reports a roughly 7-month doubling in the length of tasks frontier systems can complete, while Stanford HAI reports a sharp jump in organizational AI and generative-AI use in 2024.

■ **Deployment is broad and increasingly automation-biased.** Gallup and Anthropic both show rapid workplace diffusion, and Anthropic reports that 77% of sampled enterprise API usage is automation-oriented rather than collaborative.

■ **Aggregate unemployment is a lagging signal.** Anthropic finds little aggregate unemployment signal so far, even while Stanford finds a 13% relative employment decline for early-career workers in the most exposed occupations.

■ **Reallocation losses can persist and spill over.** BLS displaced-worker data already show incomplete reemployment after ordinary displacement; newer research shows losses are highly skewed, job-destruction spillovers depress nearby workers, and faster sectoral shifts can thicken the tail of losses.

What changes in a tail-risk framing

Conventional modeling assumption	Tail-risk adjustment used here
Past shocks eventually create offsetting labor demand in new roles.	AI can raise output with weak labor intensity and can compress multiple digital occupations at once.
Workers can retrain once and move to safer jobs.	The frontier continues to advance, so newly reskilled workers can be displaced again within a few years.
Unemployment captures most of the social burden.	Underemployment and labor-force detachment can exceed the official unemployed count.
Spillovers are local and temporary.	Displaced workers increase labor supply in adjacent human tasks, slowing wage growth and lowering job quality more broadly.

This report therefore treats the next 5-10 years as a transition window in which AI capability, adoption, and worker reallocation can fall badly out of sync.

Stress-test design

Reference path. The no-tail-shock path uses current BLS labor-force structure as the starting point and BLS 2024-34 participation projections as the structural reference. It is not a separate economic forecast; it is the benchmark the stress cases are measured against.

1. Place 78% of reference employment inside a materially exposed envelope, consistent with broad task exposure evidence in advanced economies.
2. Apply annual first-displacement waves that peak in the early-to-mid 2030s as actual AI task coverage converges toward theoretical feasibility.
3. Route displaced workers into four states: stable reemployment, underemployment, unemployment, or labor-force detachment.
4. Allow previously reemployed workers to be displaced again as the AI frontier expands (the skills treadmill effect).
5. Translate excess labor supply into wage gaps across low-, mid-, and high-barrier human task bands using labor-demand elasticity evidence as the anchor parameter and allowing more severe compression in lower-barrier work.

Scenario assumptions used to size the tail

Assumption	Severe	Extreme
Peak annual first-displacement rate (% of materially exposed employment)	4.8%	6.8%
Peak redisplacement rate (% of previously reemployed workers)	2.4%	4.4%
Trough same-year stable reemployment share	27%	18%
Peak long-term share of AI unemployment	57%	70%
Share of stable reemployed workers that still exert wage pressure on lower-tier work	10%	15%
Model status. This is an analyst-built policy stress test. It is designed to bound tail outcomes and size public buffers, not to imply false precision about exact monthly labor-market prints.		

How the labor-market rates move

The official unemployment rate rises materially in both tail cases, but the deeper signal is the interaction between unemployment and labor-force withdrawal. The reference path stays close to the demographic BLS baseline; the stress cases do not.

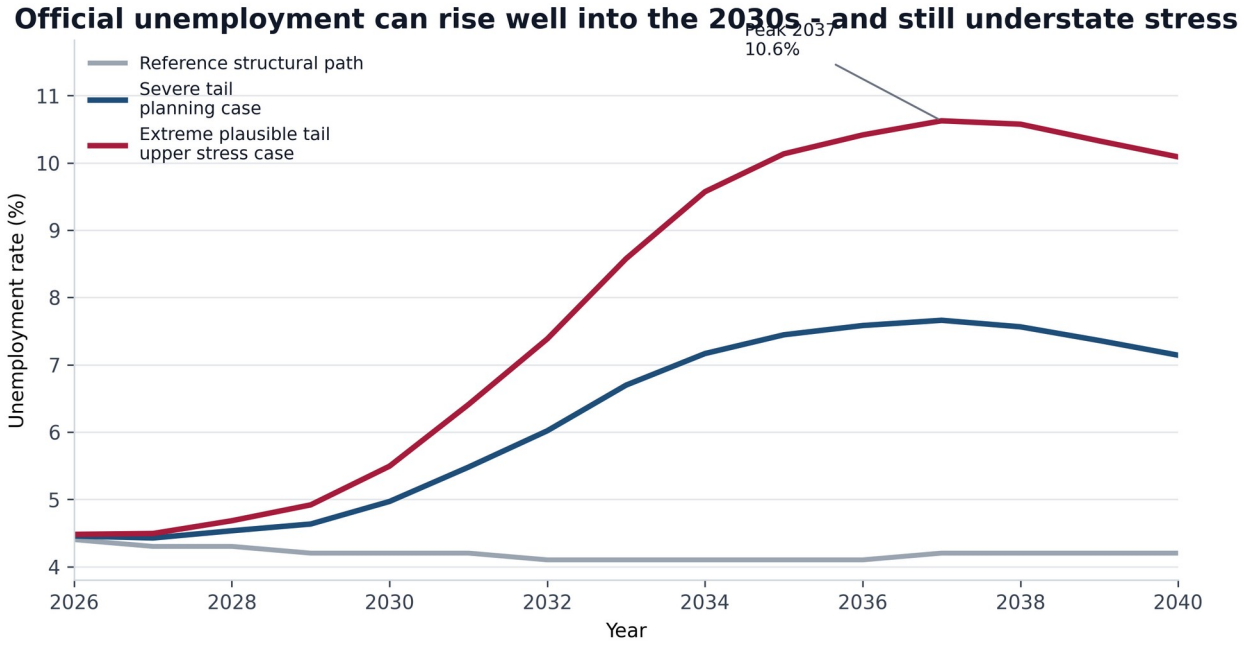

Figure 1. Official unemployment paths. Peak unemployment occurs around 2037 in both tail scenarios.

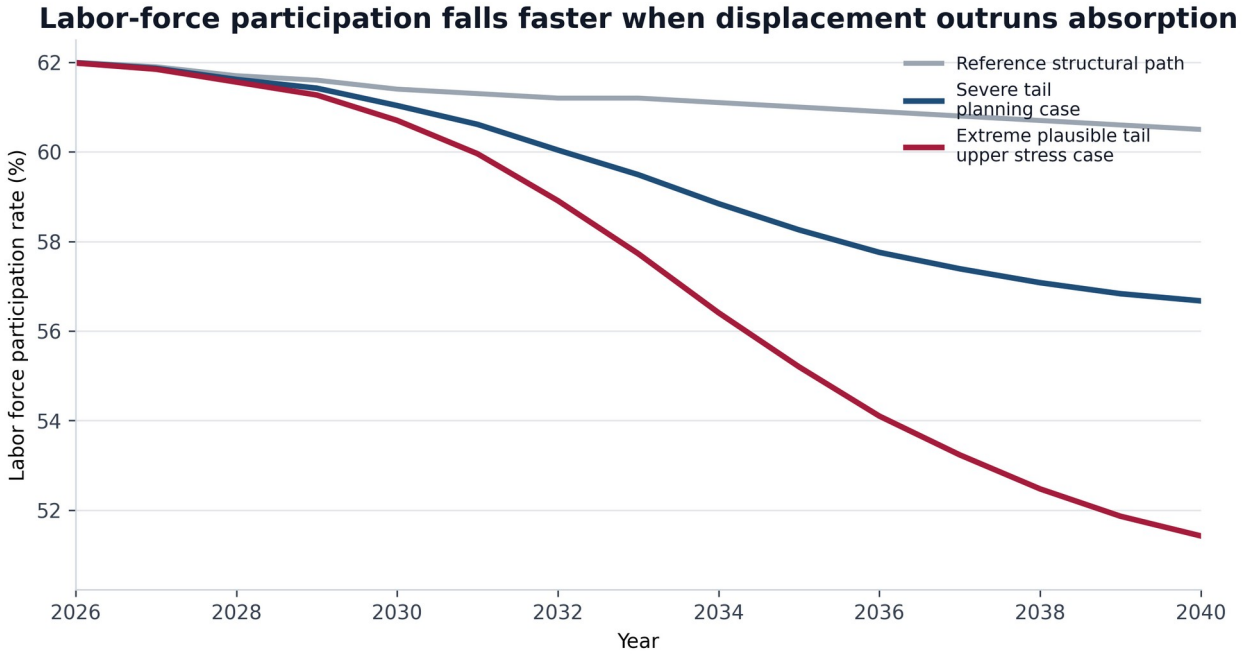

Figure 2. Labor-force participation paths. In the stress cases, official unemployment remains an incomplete measure because a larger share of displaced workers leaves the labor force.

Support-population risk

For program planning, the core metric is the AI-linked support population: unemployed + underemployed + detached / discouraged. By construction, that burden is broader than the official unemployed count and is the more relevant sizing input for safety-net design.

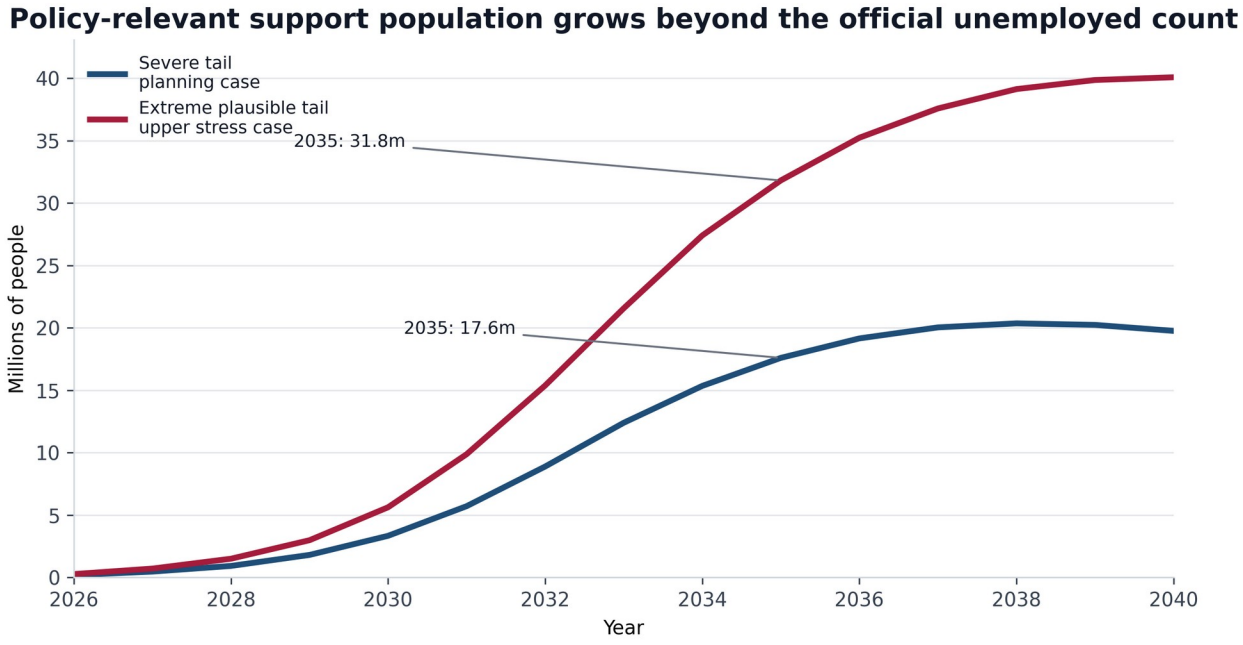

Figure 3. Total support population. In the severe case the burden peaks just above 20 million; in the extreme plausible case it continues rising into 2040.

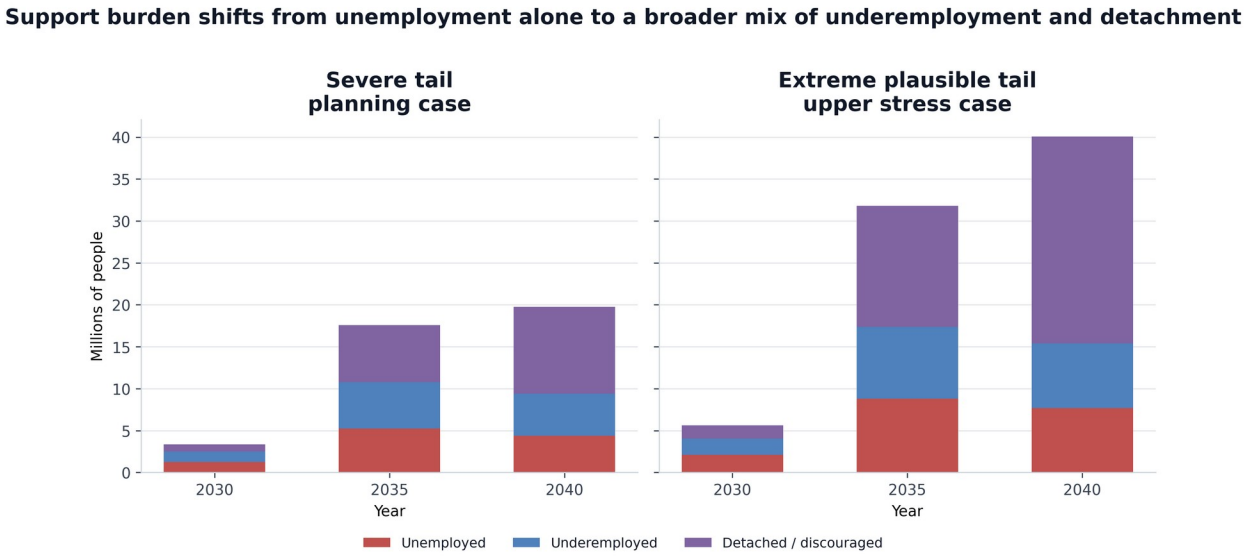

Figure 4. Composition of support burden. The tail cases evolve from mostly unemployment early on to a larger mix of underemployment and detachment.

Broader underutilization and wage compression

Underemployment matters because displaced workers can accept worse jobs without disappearing from employment statistics. Wage compression matters because those workers also increase labor supply in the tasks that remain human-intensive, especially in lower-barrier service and support work.

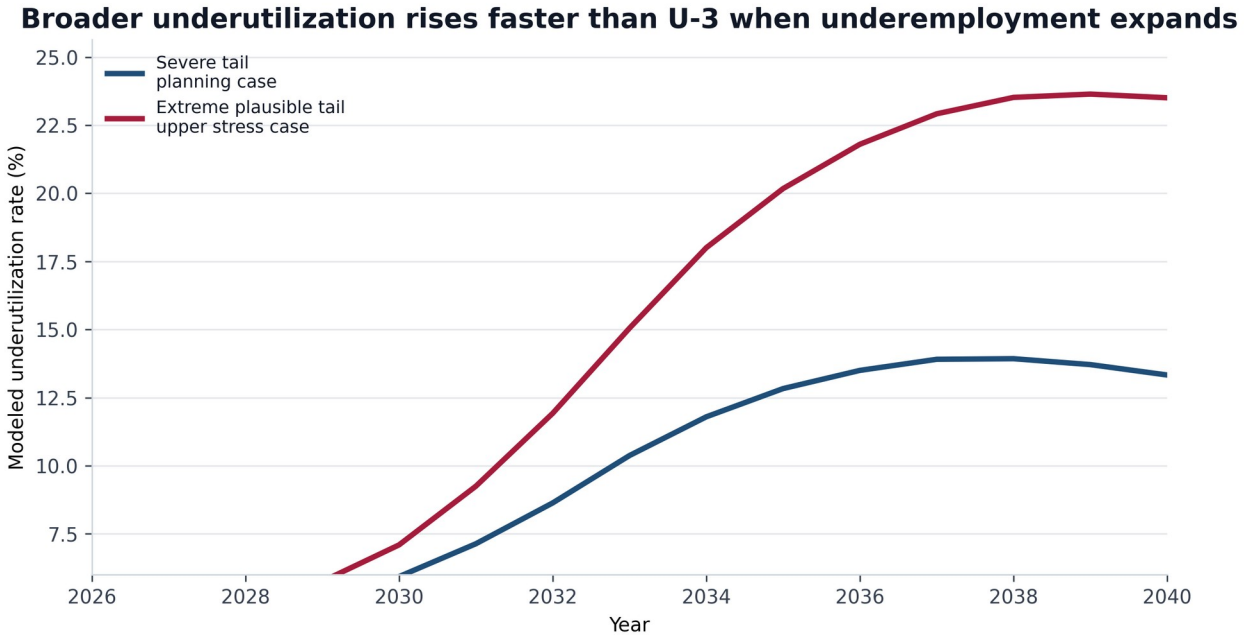

Figure 5. Modeled underutilization path. This is not the official BLS U-6; it is a stress-test indicator that adds AI-linked underemployment and partial weight on detachment.

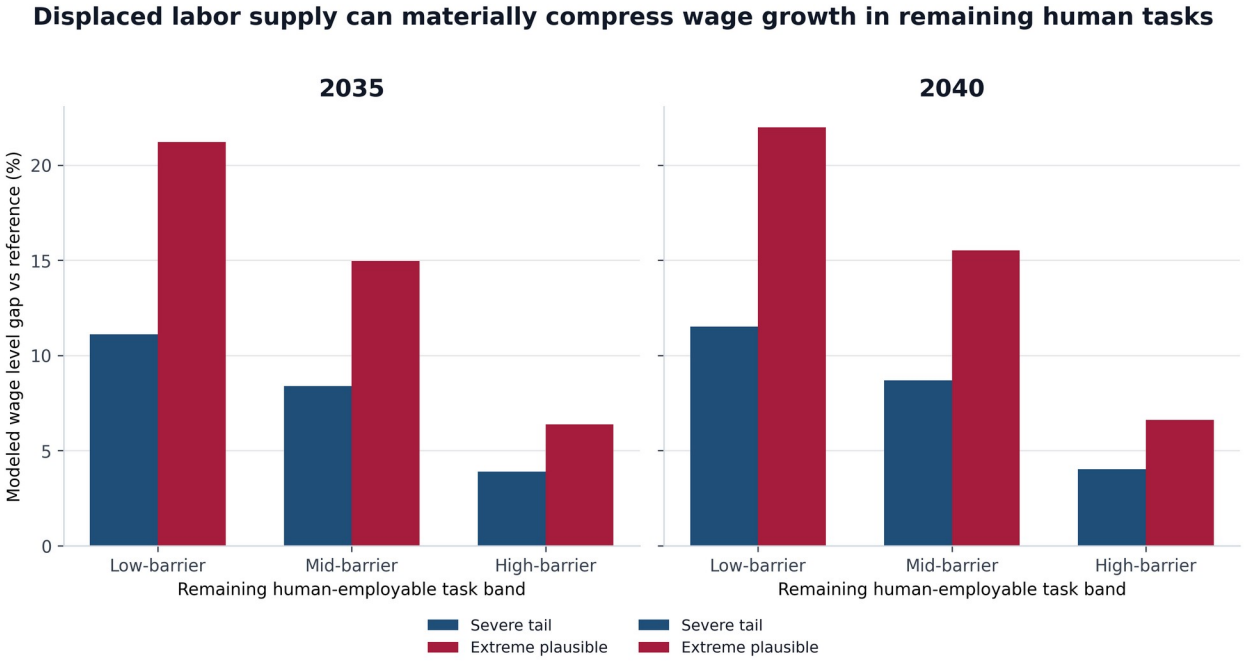

Figure 6. Cumulative wage level gaps versus the reference path. Gaps are strongest in lower-barrier human tasks because displaced labor crowds into the easiest remaining work.

Program-sizing implications

Planning implication 1. The agency should size for a support population, not a single unemployment rate. In the severe case, the 2035 burden is roughly 17.6 million and the 2040 burden is roughly 19.8 million. In the extreme plausible case, the burden is about 31.8 million in 2035 and about 40.1 million in 2040.

Planning implication 2. Long-duration joblessness can become a major subpopulation. The model yields about 2.9 million long-term AI unemployed in the severe case by 2035 and about 6.0 million in the extreme plausible case.

Planning implication 3. Wage-support design matters. Peak annual wage-growth drag occurs around 2033 and reaches about 2.2 / 1.6 / 0.8 percentage points in the severe case and about 3.9 / 2.8 / 1.2 points in the extreme case across low-, mid-, and high-barrier task bands.

Selected planning checkpoints

Year	Scenario	U-3	LFPR	Unemp.	Underemp.	Detached	Support	Low-gap
2030	**Severe**	5.0%	61.0%	1.3m	1.2m	0.9m	3.3m	-3%
2030	**Extreme**	5.5%	60.7%	2.1m	1.9m	1.6m	5.6m	-5%
2035	**Severe**	7.4%	58.3%	5.2m	5.5m	6.8m	17.6m	-11%
2035	**Extreme**	10.1%	55.2%	8.8m	8.6m	14.4m	31.8m	-21%
2040	**Severe**	7.1%	56.7%	4.4m	5.0m	10.4m	19.8m	-12%
2040	**Extreme**	10.1%	51.4%	7.7m	7.7m	24.7m	40.1m	-22%

Policy design implication. A credible response portfolio likely needs four lanes: unemployment insurance and job-search support; wage insurance / earnings supplements for downskilled reemployment; retraining paired with actual hiring channels; and reentry / social-support programs for the detached population that will not show up in headline unemployment.

Appendix A - year-by-year core planning metrics

Condensed view of the two tail scenarios. Figures are scenario outputs from this report and are shown in percent or millions of people as noted.

Year	U-3 Sev.	LFPR Sev.	Support Sev.	U-3 Ext.	LFPR Ext.	Support Ext.
2026	4.5%	62.0%	0.2m	4.5%	62.0%	0.3m
2027	4.4%	61.9%	0.5m	4.5%	61.8%	0.7m
2028	4.5%	61.6%	0.9m	4.7%	61.6%	1.5m
2029	4.6%	61.4%	1.8m	4.9%	61.3%	3.0m
2030	5.0%	61.0%	3.3m	5.5%	60.7%	5.6m
2031	5.5%	60.6%	5.7m	6.4%	60.0%	9.9m
2032	6.0%	60.0%	8.9m	7.4%	58.9%	15.4m
2033	6.7%	59.5%	12.4m	8.6%	57.7%	21.6m
2034	7.2%	58.8%	15.3m	9.6%	56.4%	27.4m
2035	7.4%	58.3%	17.6m	10.1%	55.2%	31.8m
2036	7.6%	57.8%	19.2m	10.4%	54.1%	35.2m
2037	7.7%	57.4%	20.0m	10.6%	53.2%	37.6m
2038	7.6%	57.1%	20.4m	10.6%	52.5%	39.1m
2039	7.4%	56.8%	20.2m	10.3%	51.9%	39.9m
2040	7.1%	56.7%	19.8m	10.1%	51.4%	40.1m

Appendix B - source notes and model boundaries

Important boundary. This report is a planning stress test built for tail-risk preparation. It is not claiming that the severe or extreme paths are the most likely outcomes. It is claiming that standard backward-looking models can miss them because they wait for visible labor-market damage before incorporating it.

Core evidence base used

- BLS Employment Situation, February 2026: labor force 170.5 million, unemployment 4.4%, participation 62.0%, 1.9 million long-term unemployed, and 4.4 million part-time for economic reasons.
- BLS Employment Projections / labor-supply tables, August 2025: participation projected to fall from 62.6% in 2024 to 61.1% in 2034; civilian labor force projected to rise from 168.1 million to 173.5 million.
- BLS displaced-worker evidence, 2024 release: 65.7% of long-tenured displaced workers were reemployed by January 2024, while 16.1% were unemployed and 18.2% were out of the labor force.
- OpenAI, GPTs are GPTs (2023): about 80% of the U.S. workforce could have at least 10% of tasks affected and 19% could have at least 50% affected.
- IMF AI-and-work analysis (2024/2025): about 60% of jobs in advanced economies are exposed to AI and roughly half of those exposed jobs may face negative labor-demand effects.
- Anthropic Economic Index, September 2025: rapid workplace diffusion, 40% of U.S. employees reporting AI use at work, and enterprise API usage that is 77% automation-oriented.
- Anthropic, Labor Market Impacts of AI, March 2026: actual observed coverage remains below theoretical capability, but higher observed exposure is associated with lower BLS projected employment growth and younger-worker hiring strain.
- Stanford HAI AI Index 2025: 78% of organizations reported AI use in 2024 and 71% reported generative AI use in at least one business function.
- METR, March 2025: frontier-model task horizon has doubled roughly every 7 months in recent years.
- Stanford SIEPR, Canaries in the Coal Mine (2025): early-career workers in the most AI-exposed occupations experienced a 13% relative employment decline after widespread generative-AI adoption.
- Blank and Maghzian (2024): each marginal job destroyed can impose roughly $17,000 in annual costs on other workers in the same labor market through spillovers.
- Hanushek et al. (2025) and Grigsby & Zorzi (2026): displacement losses are highly skewed and faster sectoral shifts can thicken the tail of losses even when averages improve.
- Popp (2023): own-wage labor-demand elasticity averages about -0.43 in the meta-analysis used here as an anchor for the wage-pressure module.

Modeling notes

- The reference path is anchored to current BLS totals and BLS structural participation projections. It is a benchmark, not a macro forecast.
- Support population = AI-attributable unemployed + underemployed + detached / discouraged workers.
- Modeled underutilization rate is a report-specific indicator and should not be read as the official BLS U-6.
- Wage gaps are cumulative deviations from the reference path, not nominal wage cuts, and task bands are analyst-defined simplifications of the remaining human-employable work.
- Because this is a tail-risk stress test, precision should be interpreted as scenario accounting precision, not point-forecast certainty.