Economic Scenarios for Transformative AI*

Anton Korinek, Charles I. Jones, Szymon Sacher, Tess Cotter, and Peter McCrory

The Anthropic Institute
Working Paper No. 2026-02
September 2026

Abstract

This paper presents a framework for assessing the economic consequences of AI between 2026 and 2030. In the model, AI automates a growing share of cognitive work, raising productivity and displacing workers who must search for jobs in other occupations. The model maps future paths of AI capabilities into implied paths for GDP, the labor share, wages, labor reallocation, and unemployment. We illustrate the framework by considering three scenarios: modest, substantial, and extreme. Under modest change, AI adds less than half a point to GDP growth by 2030 and raises unemployment by a tenth of a point. In the extreme change scenario, AI has transformative effects, with AI performing almost half of today's cognitive work by 2030. GDP growth then rises to 15 percent per year, the labor share of income falls from 60 to 45 percent, and nearly one in five cognitive workers is unemployed. We also surveyed US adults about their expectations for AI. Views vary widely, but the median respondent's answers are consistent with our substantial change scenario in which, by 2030, GDP rises by 8 percent and cognitive employment declines by 4 percent. The model offers a structured way to compare *possibilities* for our economic future under different expectations about AI.

*We would like to thank Daron Acemoglu, Lukas Althoff, David Autor, Tom Cunningham, Lukas Freund, Joe Hazell, Johannes Hermle, Zoë Hitzig, Ben Jones, Pete Klenow, Danial Lashkari, Eva Lyubich, Kurt Mitman, Ben Moll, Emi Nakamura, Tomas Philipson, Pascual Restrepo, David Romer, Jón Steinsson, Chris Tonetti, Ludo Visschers, and David Wiczer for insightful comments on an earlier version of this paper. We used Claude as a research and writing assistant. The views expressed in this paper are those of the authors and do not necessarily represent the views of Anthropic or The Anthropic Institute. Any remaining errors are ours.

Contents

Suggested citation for this paper: Korinek, Anton, Charles I. Jones, Szymon Sacher, Tess Cotter, and Peter McCrory. 2026. “Economic Scenarios for Transformative AI.” The Anthropic Institute Working Paper No. 2026-02.

1. Introduction

AI is likely to reshape the U.S. and global economies in profound ways in the coming decade, but how, and by how much, is extraordinarily uncertain. Published estimates of AI's future impact on productivity span orders of magnitude, and forecasts of its effect on employment range from negligible to catastrophic. Many of the estimates are based on differing analytical frameworks, which makes the projections hard to compare.

This paper takes a different approach. We build a simple, integrated economic framework that converts a small set of parameters — chiefly, what fraction of tasks in the economy are affected by AI, how widely AI is used, how large the productivity gains per task are, and how much AI automates versus augments labor — into paths for productivity, growth, wages, the labor share, job reallocation, and unemployment from now until 2030. This allows us to trace how differences in expectations about developments in AI translate into differences in economic outcomes. Our contribution is thus a structured way to compare *possibilities* for our economic future. The scenarios are not predictions, and we attach no probabilities to them; their purpose is to make the consequences of different assumptions comparable. We implement the model in a scenario explorer available at anthropic.com/institute/econ-scenarios.

The framework is based on a standard task-based model of the economy in which AI can automate or augment a growing fraction of the tasks performed by labor. Workers are employed either in cognitive occupations (management, professional, sales, and office jobs), whose work may be directly affected by AI, or in all other occupations, which are not directly exposed (construction workers or electricians, for example). AI adoption raises productivity but also displaces workers, lowering demand for labor and wages in cognitive occupations and implying that some cognitive workers must find jobs in other occupations. However, this transition is subject to frictions — for example, a software engineer may find it difficult to become an electrician — that may give rise to sustained increases in unemployment. In contrast, AI use implies higher demand for capital, which raises its return and share of total factor income. Finally, AI also speeds up the discovery of new ideas, a second channel through which it raises economic growth, though these effects are relatively small.

We parameterize three scenarios: modest, substantial, and extreme change, which broadly correspond to the wide range of published predictions. In the *modest change* scenario, GDP in 2030 is 1.6 percent higher than without AI, an effect only slightly higher than the estimate of Acemoglu (2025) and of the same order as the OECD figures of Filippucci, Gal, and Schief (2024, 2025). The *substantial change* scenario, with GDP 8.3 percent higher by 2030, corresponds to the forecasts that circulated from financial and consulting institutions in 2023 – roughly 1.5 percentage points of annual labor-productivity growth in Briggs and Kodnani (2023), and comparable magnitudes in McKinsey Global Institute (2023), Baily, Brynjolfsson, and Korinek (2023), Aghion and Bunel (2024), and the Penn Wharton Budget Model (Arnon and Smetters, 2025). The *extreme change* scenario, with GDP 32 percent higher by 2030, lies well beyond these and is best benchmarked against the literature on AGI-driven growth accelerations: the scenario analyses of Korinek (2023) and Korinek and Suh (2024), the compute-centric take-off model of Davidson (2023), the integrated assessment model of Erdil et al. (2025), and narrative accounts such as AI 2027 (Kokotajlo et al., 2025).

In each scenario, employment shifts from the cognitive occupations toward all other occupations, and cognitive wages grow more slowly than they would without AI while wages elsewhere grow faster. In the modest scenario these differences are small, even by 2030. In the extreme scenario, the effects are very large: the cognitive wage is 11.5 percent below its no-AI path while the wage in all other occupations is 34 percent above it. How much unemployment such reallocation generates depends on how many are affected by AI and how easy it is for workers to switch occupations. Across scenarios, almost all of the divergence comes after 2027. The labor share of income falls across the three scenarios as wages lag behind the increases in output. In the extreme scenario, the labor share of income falls from 60 percent today to 45 percent in 2030.

One surprising finding of our model is that the growth speedup from AI in innovation is relatively minor, even in the extreme change scenario. Our model follows the tradition of semi-endogenous growth models of Jones (1995). Automating research increases the amount of research inputs that go into the production of new knowledge, but research remains bottlenecked by physical tasks. Labor productivity rises through this channel by well under one percent in all three scenarios, and measured TFP by less. Our model, focusing on changes by 2030, also omits some important forms of feedback between research and automation. As a result, the model cannot generate the singularities of Aghion, Jones, and Jones (2019) or the explosive research-driven acceleration of Davidson, Halperin, Houlden, and Korinek (2026), so the innovation effects here are likely a lower bound.

To complement the three scenarios, we also conducted a representative survey of the US population. We asked 10,980 US adults: (a) when AI will be able to do each of eight different cognitive tasks of increasing difficulty, (b) how widely AI will be used, (c) how large the productivity gains will be, (d) whether workers will be automated or augmented, and (e) how long a displaced worker will take to find a new job. Running the simulations using their median answers as inputs to the model, we find that the public's expectations land close to our substantial change scenario.

The labor market is where the stakes of our economic scenarios are highest. In the macroeconomic data to date, the labor market effects of AI have been relatively muted even as measured AI usage has grown rapidly (though see Humlum and Vestergaard, 2025; Azar et al., 2025, on starting wages in exposed occupations where evidence is mixed). Two interpretations are consistent with this fact. One is that we are in a modest change scenario and the economic effects will remain small. The other is that we are only at the very beginning of a more extreme scenario that will reshape the role and earnings of labor. The two interpretations have very different policy implications. If the effects of AI on the economy are modest, the adjustments required of existing institutions are correspondingly small. If instead AI substantially raises output while displacing work, the mechanisms by which people receive resources from an immensely more prosperous economy may themselves need to change. The range of reasonable options, from worker retraining to new forms of insurance and income support, is wide. The framework can represent either case and shows what would have to be true of AI in each.

Caveats. Like every model, our framework is a stark simplification of a complex reality: it isolates a few key forces and omits many others that may become relevant and important in the coming years. Several caveats should be made explicit. First, our framework omits many potentially relevant forces, for example, catastrophic risks, political economy considerations, business cycles, and possible financial market disruptions. All of these could be important in any of our scenarios.

Other omissions may be examined in future extensions to this framework. We do not consider price rigidities and the associated demand effects (beyond the slow adjustment of real wages), so the model cannot generate the negative feedback in which disruption depresses demand and amplifies its own labor-market consequences. We do not assign a value to AI's use in household production, which lies largely outside GDP (Nguyen et al., 2026). Worker heterogeneity is very coarse: workers differ only in employment status and in which of two groups of occupations they belong to; within a group all earn the same wage. A displaced worker who moves to the other group earns that group's wage at once, with no discount for lost tenure or occupation-specific skills (Jacobson, LaLonde, and Sullivan, 1993; Huckfeldt, 2022; Braxton and Taska, 2023). In future versions, it may be desirable to include heterogeneity by tenure and skill, across firms by adoption capacity, across regions by industrial composition, and across demographic groups in mobility and retraining costs, which would permit a richer discussion of inequality.

The factor structure is also deliberately simple: labor and a single capital good supplied along an exogenous schedule with fixed elasticity, which does not distinguish compute from other capital and does not explicitly model saving decisions. While we allow AI to improve the productivity of newly automated tasks over time, we do not currently capture productivity growth of tasks already performed by capital before 2024. This omits a potentially important channel for both growth and for sustaining the labor share.

We also consider only the effects of AI on cognitive tasks, and do not allow for rapid advances in robotics and their effects on physical tasks in the coming years. Largely for this reason, we do not extend the analysis beyond 2030.

Literature. This paper draws on two strands of literature. Our baseline model builds on the task-based automation literature: the formal structure of Zeira (1998), the characterization of occupations by their task content in Autor, Levy, and Murnane (2003) and Acemoglu and Autor (2011), and the displacement-reinstatement decomposition of Acemoglu and Restrepo (2019) and Autor and Salomons (2018). Autor, Chin, Salomons, and Seegmiller (2024) measure employment in new work. Our model also builds on the analysis of worker-replacing technical change in Korinek and Stiglitz (2019) and the comprehensive framework of growth under transformative AI in Trammell and Korinek (2026). We model the supply of capital following the long-run wealth schedule of Moll, Rachel, and Restrepo (2022).The model of innovation follows the semi-endogenous growth literature (Romer, 1990; Jones, 1995) and the associated empirical findings of Bloom, Jones, Van Reenen, and Webb (2020).

The labor market block draws on the search-and-matching tradition of Diamond (1982), Mortensen and Pissarides (1994), and Pissarides (2000). AI displaces workers in cognitive occupations who then search for new work, primarily in the rest of the economy. We adopt the matching function introduced by den Haan, Ramey, and Watson (2000) that determines hiring as a function of generalized labor market tightness that accommodates lower search intensities when switching occupations (see Abraham, Haltiwanger, and Rendell (2020)). Switching occupations can be hard due to occupation-specific human capital, as documented in Kambourov and Manovskii (2009) . We model this with a discount on cross-occupation search effort among unemployed workers. Layoffs arise due to real wage rigidities a la Blanchard and Galí (2007) and job rationing in the sense of Michaillat (2012).

A recent literature emphasizes that occupations are bundles of tasks, and workers differ in their productivity across them (Autor and Thompson, 2025; Freund and Mann, 2026; Althoff and Reichardt, 2026). We do not currently incorporate this richness. These models generate earnings losses for displaced workers and wage gains or losses within an exposed occupation. In Autor and Thompson (2025), automating the inexpert tasks of an occupation raises the expertise and the wage of the job that remains, while automating the expert tasks lowers both. In Fukui, Nakamura, and Steinsson (2026), technology that standardizes tasks lowers wages without lowering employment.

Outline. Section 2 lays out the model to study how AI automates the production of goods and services, how it raises research input and thereby growth, and how frictional reallocation of workers creates unemployment. Section 3 calibrates the parameters, specifies the three scenarios, and lays out our survey methodology. Section 4 presents the results of the scenarios we defined and of what US adults expect. Section 5 concludes. Appendix A describes the simulation strategy, Appendix B details the survey of US adults, and Appendix C provides derivations for the innovation block of our model.

2. Model

In a nutshell. Improvements in AI raise productivity and, over time, make the economy richer. These gains are distributed in two ways: through higher returns to the owners of capital and through changes in wages. In the long run, in standard neoclassical models, capital is perfectly elastic, since new machines can always be built. All the productivity gains then accrue to labor and raise wages, while capital earns its usual (unchanged) long-run return.

However, in the short and medium run — which includes the time between today and 2030 in our framework — capital is less elastic and slower to adjust, so some of the productivity gains show up as a higher return to the owners of capital. How the gains are split between capital and wages depends on how easily the economy can add capital.

The second distributional issue is how wages differ between the cognitive occupations and all other occupations. In the long run, after all the adjustments have occurred, there is a single wage in the model, and it typically rises as a result of AI. But in the medium run, AI reduces the demand for cognitive labor by automating its tasks. This reduction shows up as lower wages and lower employment in cognitive occupations (relative to the paths they would have taken without AI). How much is a decline in relative wages and how much a decline in employment depends on another element in the model, the extent of real wage rigidity. If firms are slow to reduce wages, relative wages decline less and the fall in demand shows up instead as lower employment in the cognitive occupations.

The workers displaced from cognitive work must switch occupations and find jobs in the rest of the economy. Because search and matching take time, this raises unemployment among the workers affected by AI.

Finally, AI also speeds up the discovery of new ideas. Ideas are produced from research in a semi-endogenous growth setup, so a one-time increase in research input raises the long-run *level* of the stock of ideas rather than its long-run growth rate. AI's uplift to research input therefore adds a little to the growth rate of ideas each year, and these additions accumulate slowly. Between now and 2030, the effect on the level of technology is typically small.

We now set out the model formally. Here are the key ingredients:

- Production involves a combination of tasks, where each task is performed by labor or by capital. The final goods market is perfectly competitive.
- Labor is supplied exogenously while capital is supplied along an upward-sloping schedule.
- There are two groups of occupations, of which only the cognitive group is exposed to AI. Frictions and a search/matching process govern unemployment.
- A *scenario* is a path for five objects that describe AI in production — the mass of tasks AI affects, its diffusion across them, the gain per task, the share of AI use that automates rather than augments, and the number of new labor tasks per automated task — together with two frictions of the labor market.

We present the model's predictions relative to a counterfactual economy without AI. The gap between a particular scenario and the counterfactual economy without AI is denoted by Δ. For example, $\Delta \ln Y_t$ is the log difference in output in an economy with AI to one without at time t.

We first discuss the impact of AI on the level of output and employment, then its impact on growth through ideas, and then its impact on unemployment.

2.1 Production and factor shares

2.1.1 Technology and factor markets

Output is produced by a constant-returns technology defined over task instances, in the task-based tradition of Zeira (1998), Autor et al. (2003), Acemoglu and Autor (2011), and Acemoglu and Restrepo (2018). Specifically, output is a CES aggregate over task instances,

$$Y_t = \left(\sum_i \omega_i^{\frac{1}{\sigma}} \, y_{i,t}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \tag{1}$$

where $y_{i,t}$ is the output of an instance of task i and the weights ω_i, with $\sum_i \omega_i = 1$, are base-period expenditure shares. A task is a type of work, such as reviewing a contract or taking a patient's history. An instance is a single performance of a task, for example one contract review.

A task consists of a great many instances, and the sum in Equation (1) runs over instances, with the instances of a task sharing its weight equally. AI arrives instance by instance: some contracts are reviewed with AI while others are not. The two types of instances are separate entries in the production function, with their own prices and quantities, rather than a single aggregated composite. The parameter σ is the elasticity of substitution across instances, of the same task or of different tasks. We refer to it as the elasticity across tasks. In particular, an AI-performed instance and a human-performed instance of the same task are imperfect substitutes with the same elasticity σ.

We set $\sigma = 0.5$, so that tasks are gross complements, the value assumed by Acemoglu and Restrepo (2022), which traces to the estimate of Humlum (2019) across three broad occupation groups. Estimates range from below one to above one (e.g. based on occupations, as in Burstein, Morales, and Vogel, 2019). Jones and Tonetti (2026) bound the task elasticity from evidence on the capital-labor elasticity of substitution and assume an even lower value.

Each task instance can be produced by labor or by capital:

$$y_{i,t} = A_t \, \alpha_{L,i,t} \, \ell_{i,t} + \alpha_{K,i,t} \, k_{i,t}, \tag{2}$$

where $\ell_{i,t}$ and $k_{i,t}$ are labor and capital per instance, $\alpha_{L,i,t}$ and $\alpha_{K,i,t}$ are task-specific factor productivities, which AI shifts over time, and A_t is the stock of ideas, which augments labor on every task and is taken as given in this section. The labor-augmenting structure for A_t facilitates a balanced growth path in our model in the absence of AI; letting it augment both labor and capital has only small effects in what follows (though it would have larger effects over longer horizons). Capital and labor are therefore perfect substitutes *within* a task and, since $\sigma < 1$, gross complements *across* tasks. Once each instance is assigned to a factor, the technology in Equations (1)–(2) is a constant-returns aggregate of total capital $K_t = \sum_i k_{i,t}$ and labor $L_t = \sum_i \ell_{i,t}$.

In the absence of unemployment, which is introduced in Section 2.3, labor L_t is supplied exogenously and is paid the wage w_t, which clears the labor market. Capital K_t is rented at the rate r_t, which clears the market for capital, whose supply we specify in a moment. Markets are competitive, so each task instance is priced at unit cost, and cost minimization assigns it to the cheaper factor,

$$p_{i,t} \;=\; c_{i,t} \;=\; \min\left\{\frac{w_t}{A_t\,\alpha_{L,i,t}},\;\frac{r_t}{\alpha_{K,i,t}}\right\}, \tag{3}$$

so that capital performs task instance i at date t if and only if $\alpha_{K,i,t}/(A_t\alpha_{L,i,t}) > r_t/w_t$. Ranking instances by capital's comparative advantage $\alpha_{K,i,t}/\alpha_{L,i,t}$, the equilibrium assignment is a threshold rule, as in Acemoglu and Restrepo (2018). With the final good as the numeraire, the CES structure in Equation (1) implies the expenditure shares

$$s_{i,t} \;\equiv\; \frac{p_{i,t}\,y_{i,t}}{P_t Y_t} \;=\; \omega_i\left(\frac{c_{i,t}}{P_t}\right)^{1-\sigma}, \qquad P_t \;=\; \left(\sum_j \omega_j\, c_{j,t}^{1-\sigma}\right)^{\frac{1}{1-\sigma}} \;\equiv\; 1, \tag{4}$$

where P_t is the price index of the final good.

Let $\mathcal{L}_t$ denote the set of task instances performed by labor at date t. The labor and capital shares and the wage identity are then

$$s_{L,t} \;=\; \sum_{i\in\mathcal{L}_t} s_{i,t} \;=\; \frac{w_t L_t}{P_t Y_t}, \qquad s_{K,t} \equiv 1 - s_{L,t}, \qquad \Delta\ln w_t \;=\; \Delta\ln(Y_t/L_t) \;+\; \Delta\ln s_{L,t}. \tag{5}$$

The wage identity holds exactly at every date, and we use it repeatedly: any two of the wage, output per worker, and the labor share determine the third. We set $s_{L,t_0} = 0.6$.

In the base period, $t_0 = 2024$, the last year before the use of large language models spread, we choose units so that all unit costs are equal and $s_{i,t_0} = \omega_i$. We measure the importance of each task by its initial share of total labor payments: $m_i \equiv s_{i,t_0}/s_{L,t_0} = \omega_i/s_{L,t_0}$. By construction, then, $\sum_i m_i = 1$, where the sum is over the labor-performed tasks. After t_0 the cost shares $s_{i,t}$ move with relative unit costs.

The supply of capital. The demand for capital follows from the capital share, $K_t = s_{K,t}\,P_t Y_t/r_t$. Its supply is an upward-sloping schedule in the rental rate: relative to the path the economy would follow without AI, on which the rental rate is constant at $\bar r$, the capital stock is higher by ε percent for every percent by which the rental rate exceeds $\bar r$,

$$\Delta\ln K_t \;=\; \varepsilon\,\Delta\ln r_t. \tag{6}$$

The schedule is the reduced form of Moll, Rachel, and Restrepo (2022), whose households accumulate savings and hold more wealth in the long run the higher the return. As $\varepsilon \to \infty$ the rental rate is pegged at $\bar r$, the case of a small economy in world capital markets, and any amount of capital is forthcoming at an unchanged return. At $\varepsilon = 0$ the capital stock does not respond to AI at all within the horizon. The returns accrue to capital's owners. We report the rental rate as a net return $r_t - \delta$, with $\bar r = 0.115$ and depreciation of $\delta = 0.05$ a year, so that the net return is 6.5 percent without AI and the capital-output ratio is $s_{K,t_0}/\bar r = 3.5$.[1]

[1] Moll, Rachel, and Restrepo write their schedule with a semi-elasticity in the net return, $\Delta\ln K_t = \eta\,\Delta(r_t - \delta)$, and set $\eta = 50$. With δ constant, $\Delta(r_t - \delta) \approx \bar r\,\Delta\ln r_t$, so the two schedules agree to first order at $\bar r$ when $\varepsilon = \eta\,\bar r$, which is $50 \times 0.115 \approx 6$. They call $\eta = 50$ conservative because much of the evidence points to a less elastic supply of capital,

The economy without AI. In the absence of AI, we assume the economy is on a balanced growth path. The stock of ideas A_t grows at rate g, and the labor force L_t at rate n. The rental rate is constant at $\bar{r}$, and every expenditure share is then constant. The wage and output per worker therefore grow at g, capital at $g+n$, and measured TFP at $s_{L,t_0}\, g$. We set $g = 1.67$ percent a year, so that measured TFP grows at $g_A \equiv s_{L,t_0}\, g = 1$ percent, and we take the labor force to grow at $n = 0.33$ percent a year, within the range of the Bureau of Labor Statistics projections. GDP without AI therefore grows at $g + n = 2$ percent a year.

The two groups of occupations. The labor-performed tasks are sorted into two groups of occupations, which we think of as islands: the *cognitive* occupations C (management, professional, sales, and office work), whose tasks AI affects, and *all other* occupations N (manual and interpersonal work, say), whose tasks it does not. We refer to them as the two *groups*, the cognitive group and the all-other group. Their tasks account for shares s_{C,t_0} and s_{N,t_0} of base-period expenditure, with $s_{C,t_0} + s_{N,t_0} = s_{L,t_0}$, and they employ $\ell_{C,t}$ and $\ell_{N,t}$ workers. We measure employment, and later the pool of workers between jobs, as shares of the labor force L_t, which grows at the rate n with or without AI, and we write L without a subscript for the labor force in these units, a constant, so that $\ell_{C,t} + \ell_{N,t} = L$. Workers are homogeneous in production. In the full-employment economy of this section, with workers free to move between the groups, they earn a common wage w_t wherever they work. During the transition of Section 2.3 each group is paid its own wage.

Two facts about these groups are consequential for the model. First, AI affects C and not N, so some workers in C have to switch occupations as AI spreads. Second, moving between groups requires search, and that raises unemployment. With the labor force fixed, what AI changes therefore is the division of workers between the two groups and the fraction of workers who are between jobs.

2.1.2 AI scenarios

An AI scenario is largely determined by a set of paths for five exogenous objects, which describe AI in the production of goods and services. At t, AI affects a subset $\mathcal{A}_t$ of the tasks in the economy, all of them cognitive tasks, and the *affected mass* is the mass of that set, $m_t \equiv \sum_{i\in\mathcal{A}_t} m_i$: the fraction of the economy's tasks that AI affects.

On each affected task i, the *diffusion share* $d_{i,t}$ is the fraction of the task's instances actually performed with AI, and the *gain* $a_{i,t}$ is the log productivity gain per AI-performed instance, or equivalently the log decline in the instance's unit cost. In the aggregate, $d_t \equiv \sum_{i\in\mathcal{A}_t} m_i\, d_{i,t}/m_t$ is the mass-weighted diffusion share and $a_t \equiv \sum_{i\in\mathcal{A}_t} m_i\, d_{i,t}\, a_{i,t} / \sum_{i\in\mathcal{A}_t} m_i\, d_{i,t}$ the average gain over AI-performed instances, so that $\sum_i m_i\, d_{i,t}\, a_{i,t} = m_t\, d_t\, a_t$ identically.

All three objects are measured cumulatively from the pre-AI benchmark. The products in which they enter the rest of the paper are read the same way as m_t: $m_t d_t$ is the fraction of the

and their economy takes decades to reach the steady states it describes. We therefore set $\varepsilon = 3$, half their value. In their units, a rise of one percentage point in the net return then raises the capital stock by about $3/0.115 \approx 26$ percent.

economy's task instances that AI performs, and $m_t d_t a_t$ is that fraction times the gain on each.

Performing a task instance with AI takes one of two forms. Under *augmentation*, AI is assumed to raise a worker's labor productivity $\alpha_{L,i,t}$ on that instance by $a_{i,t}$ log points. Under *automation*, an AI system performs the instance outright, so that the instance is performed by capital. In this case we assume that AI raises capital's productivity $\alpha_{K,i,t}$ to the point where capital, rented at the no-AI rate $\bar{r}$, supplies the instance at unit cost $\bar{r}/\alpha_{K,i,t} = e^{-a_{i,t}}\,\bar{w}_t/(\bar{A}_t\alpha_{L,i})$, that is, $a_{i,t}$ log points below labor's unit cost without AI, where bars mark the no-AI path. In either form, the unit cost of the instance falls by $a_{i,t}$ at no-AI factor prices.

The fourth scenario primitive, the *automation share* $\psi_t \in [0,1]$, is the fraction of AI-performed instances that are automated rather than augmented. The technology determines whether an instance is automated or augmented, and firms take this as given. The gain on an automated instance is available only if capital performs it. The gain on an augmented instance is available only if a worker performs it.[2]

The fifth scenario object, the *reinstatement ratio* $\rho \in [0,1]$, is the mass of newly created labor tasks as a fraction of the mass of automated tasks, in the spirit of Acemoglu and Restrepo (2019) and Autor et al. (2024). That is, a mass $\psi_t\, m_t\, d_t$ of existing labor tasks is automated and a mass $\rho\, \psi_t\, m_t\, d_t$ of new labor tasks is created, so the mass of labor tasks changes by $-(1-\rho)\,\psi_t\, m_t\, d_t$ on net. The new tasks belong to the cognitive occupations, whose tasks were automated, and enter at the same unit cost as the remaining labor tasks. Reinstatement therefore offsets displacement one-for-one in the affected wage bill, adds no separate variety or weak-link term to output, and leaves the factor-price frontier unchanged. Labor is mobile across the tasks of a group, so a new task is staffed from the moment it enters. We follow the approach of Acemoglu and Restrepo (2018) in which new labor-intensive tasks enter at one end of a fixed range of tasks while old capital-performed tasks drop out at the other.

The affected mass and the diffusion share follow logistic curves, and the log gain follows a linear path:

$$m_t = \frac{\bar{m}}{1+e^{-\kappa_m\,(t-t_m)}}, \qquad d_t = \frac{\bar{d}}{1+e^{-\kappa_d\,(t-t_d)}}, \qquad a_t = a_0 + g_a(t-t_0). \tag{8}$$

The automation share ψ_t and the reinstatement ratio ρ are held constant in the three scenarios.

Throughout the remainder of this paper, we report exact solutions to our model equations with equality = and first-order approximations with the symbol $\approx$, which often give better intuition. The tables in Appendix A give the full recipe for solving the model.

2.1.3 Measured TFP and the factor-price frontier

On the fraction $d_{i,t}$ of task i's instances that are performed with AI, unit cost falls by $a_{i,t}$ log points, whether a worker performs the instance with AI or capital performs it outright. Following

[2] In principle, a firm could still hand an automated instance back to a worker without AI. By Equation (3) it keeps the instance on capital as long as

$$\Delta \ln r_t \;-\; \Delta \ln w_t \;\le\; a_{i,t}, \tag{7}$$

that is, as long as AI has not raised the rental rate relative to the wage by more than the gain itself.

Acemoglu (2025), a growing literature uses Hulten's theorem (Hulten, 1978) to compute the resulting gain in aggregate productivity as the sum of these cost declines, each weighted by its Domar weight, which in our model is the task's expenditure share ω_i. Dividing the Domar weight ω_i by s_{L,t_0} gives the mass m_i, so measured TFP rises by

$$\Delta \ln \mathrm{TFP}_t \;\approx\; \sum_i \omega_i\, d_{i,t}\, a_{i,t} \;=\; s_{L,t_0} \sum_i m_i\, d_{i,t}\, a_{i,t} \;=\; s_{L,t_0}\, m_t\, d_t\, a_t. \tag{9}$$

This expression does not involve the automation share since the impact on aggregate productivity depends only on the decline in an instance's unit cost, which is $a_{i,t}$ whether AI is used in an automated or augmented way.

Whether the gain accrues to workers or capital owners is a separate question. With constant returns and competitive pricing, factor payments exhaust output: $P_t Y_t = w_t L_t + r_t K_t$. Measured TFP growth equals its dual, the share-weighted growth of the two factor prices. With the final good as numeraire, $d \ln \mathrm{TFP}_t \;=\; s_{L,t}\, d \ln w_t \,+\, s_{K,t}\, d \ln r_t$. Applied to the productivity gain in Equation (9),

$$s_{L,t_0}\, \Delta \ln w_t \;+\; s_{K,t_0}\, \Delta \ln r_t \;\approx\; s_{L,t_0}\, m_t\, d_t\, a_t. \tag{10}$$

That is, between them the two factor prices absorb the entire TFP gain, whatever the values of ψ_t, ρ, and σ. If the rental rate does not move, as at $\varepsilon = \infty$, the wage rises by the full TFP gain, $\Delta \ln w_t \approx m_t d_t a_t$ (Caselli and Manning, 2019).

If capital is slow to adjust and the rental rate rises, then the wage will rise by less in response to an increase in TFP. How much an increase in rental rates reduces the change in wages is given by (10) . Each percent increase in rental rates comes with a $s_{K,t_0}/s_{L,t_0} = 2/3$ of a percent lower wage. If capital is sufficiently inelastic, the wage can actually fall relative to its no-AI path, as would be the case when $s_{K,t_0} \Delta \ln r_t > s_{L,t_0} m_t d_t a_t$.

Equations (9) and (10) hold to first order, but a benefit of Hulten's theorem is that they do so for any constant-returns technology in an efficient economy. The gains from AI may be large so a first-order approximation may not be accurate enough for the more extreme scenarios we consider. We use exact solutions in our simulations but describe some of the intuition using first-order approximations.

To illustrate, take the parameter values that we will assume for the substantial change scenario at the start of 2030: AI affects $m_t = 30$ percent of the economy's tasks, it is actually used on $d_t = 40$ percent of their instances, and it lowers the unit cost of each such instance by $a_t = 0.45$ log points. AI then performs $m_t d_t = 12$ percent of the economy's task instances, and by Equation (9), AI raises measured TFP by

$$\Delta \ln \mathrm{TFP}_{2030} \;\approx\; 0.6 \times 0.30 \times 0.40 \times 0.45 \;=\; 0.032,$$

about 3 percent. (The exact solution is 0.029.) If the rental rate stayed at $\bar{r}$, as it would once capital has fully adjusted, the factor price frontier (10) implies that the whole gain accrues to labor, and the wage would be $m_t d_t a_t = 0.054$ above its no-AI path. When the supply of capital is less than perfectly elastic, the rental rate rises and both factors share the gains: at $\varepsilon = 3$ the rental rate in 2030 is 4.6 percent above $\bar{r}$, so capital's term in Equation (10) is $0.4 \times 0.045 = 0.018$, and the wage is 1.9 percent above its no-AI path rather than 5.5 percent.

2.1.4 First-order solutions

The first order solutions for the model are then summarized in the following key equations:

$$\begin{aligned}
\Delta \ln w_t &\approx m_t\, d_t\, a_t \;-\; \tfrac{s_{K,t_0}}{s_{L,t_0}}\, \Delta \ln r_t, \\
\Delta \ln r_t &\approx \frac{1}{\varepsilon + \sigma / s_{L,t_0}} \left[m_t\, d_t\, a_t \;+\; \big((1-\rho) - (1-\sigma)\, a_t\big)\, \frac{\psi_t\, m_t\, d_t}{s_{K,t_0}} \right], \\
\Delta \ln s_{L,t} &\approx -\,(1-\rho)\, \psi_t\, m_t\, d_t \;+\; (1-\sigma)\, \psi_t\, m_t\, d_t\, a_t \;-\; (1-\sigma)\, \tfrac{s_{K,t_0}}{s_{L,t_0}}\, \Delta \ln r_t, \\
\Delta \ln(Y_t/L_t) &\approx m_t\, d_t\, a_t \;+\; (1-\rho)\, \psi_t\, m_t\, d_t \;-\; (1-\sigma)\, \psi_t\, m_t\, d_t\, a_t \;-\; \sigma\, \tfrac{s_{K,t_0}}{s_{L,t_0}}\, \Delta \ln r_t.
\end{aligned} \tag{11}$$

The wage equation derives from the TFP growth expression just discussed in equation (10).

Equating the capital demanded to the capital supplied along the schedule (6) gives the rental-rate gap. The bracket in the $\Delta \ln r_t$ equation is the capital that AI calls for at an unchanged rental price: capital keeps pace with the output that the gains themselves produce, $m_t d_t a_t$. Further, each automated instance shifts its wage bill to capital and raises output by the same amount, which together call for $\psi_t m_t d_t / s_{K,t_0}$ more capital (less reinstatement and less the expenditure that substitution sends back toward labor's tasks). The rental rate rises by enough to reconcile that demand with the supply schedule. Demand falls with the rental rate at the elasticity $\sigma/s_{L,t_0}$ and supply rises with it at ε.

The labor share responds through three channels. First, automation reassigns tasks from labor to capital. The entire wage bill of an automated instance moves to capital, however small the cost saving that motivated the switch. Net of reinstatement, then, the labor share falls by $(1-\rho)\, \psi_t\, m_t\, d_t$. Second, the cost declines on the instances that labor keeps offset only $(1-\psi_t)\, m_t\, d_t\, a_t$ of the rise in the wage. Relative to the automated instances, the remaining labor instances become more expensive, by $\psi_t\, m_t\, d_t\, a_t$ on average. Because tasks are gross complements ($\sigma < 1$), demand substitutes away from these instances less than in proportion to their price. They are the weak link, so their expenditure share rises, which adds back $(1-\sigma)\, \psi_t\, m_t\, d_t\, a_t$. Third, by the same complementarity, when the rental rate rises, capital's instances become dearer relative to labor's and expenditure follows them, which lowers the labor share by $(1-\sigma)(s_{K,t_0}/s_{L,t_0})\, \Delta \ln r_t$.

Output per worker can then be obtained from the labor share in (5). And capital similarly is derived from the capital share, $\Delta \ln K_t = \Delta \ln s_{K,t} + \Delta \ln(Y_t/L_t) - \Delta \ln r_t$.

Two insights follow directly. First, at a fixed rental rate, the wage tracks productivity, $m_t\, d_t\, a_t$, not net displacement, $(1-\rho)\, \psi_t\, m_t\, d_t$. Displacement captures the change in the number of tasks that labor performs. It moves the labor share and the demand for capital, and it affects the wage only through the rental rate. When capital is less than perfectly elastic, its rental price rises and the wage is affected. This explains the difference vis-a-vis Caselli and Manning (2019), who show that when capital is perfectly elastic (e.g. in the long run), the interest rate is fixed and technical change cannot lower the average wage.[3]

This leads directly to the second insight: under what conditions can the wage decline? The answer depends on the displacement and the cost saving. Automation transfers to capital the wage bill of the tasks it takes over, $(1-\rho)\psi_t m_t d_t$, whether the cost saving on those tasks is large or small. In contrast, the cost saving from (10), $s_{L,t_0} m_t d_t a_t$, is the gain available to distribute. When the saving per automated task is small relative to the transfer and when little capital is forthcoming, the rental price has to rise by enough to rationalize the transfer. If this rise is more than the cost saving, the wage is below its no-AI path: labor has lost the tasks and, instead of working with cheap and plentiful capital, it is limited by capital that is scarce and expensive. This is the "so-so automation" of Acemoglu and Restrepo (2019) in which displacement outweighs a modest productivity gain. In 2030 at our baseline $\varepsilon = 3$, the wage rises in every scenario. However, the wage falls if the capital supply elasticity ε is sufficiently small, as we will see.

2.1.5 Employment in the two groups

How many workers need to reallocate from cognitive occupations to all other occupations? Write $\ell^*_{o,t}$ for employment in group o when workers move freely between groups at no cost and all workers are paid a common wage. We call this allocation the groups' *target*, since it is what employment is adjusting toward. Interestingly, the target for cognitive occupations is pinned down by looking at all other workers unaffected directly by AI.

> **Key result.** AI does not affect the productivity of tasks of all other occupations, N, so the cost of each depends on the wage, not the automation gain. The demand for labor on those tasks is the ordinary CES demand for a good of price w_t: it rises one for one with output and falls with the wage, with elasticity σ. In other words, as long as AI raises the wage along with output, N's employment rises by *less* than output. Because labor supply is fixed, the workers N gains are the workers C loses, so C's proportional decline is N's proportional rise scaled by the relative size of the two groups, $s_{N,t_0}/s_{C,t_0}$:
>
> $$\tilde{\ell}_{N,t} \;\equiv\; \ln\frac{\ell^*_{N,t}}{\ell_{N,t_0}} \;=\; \Delta\ln(Y_t/L_t) \;-\; \sigma\,\Delta\ln w_t, \qquad \frac{\ell_{C,t_0}-\ell^*_{C,t}}{\ell_{C,t_0}} \;=\; \frac{s_{N,t_0}}{s_{C,t_0}}\,\frac{\ell^*_{N,t}-\ell_{N,t_0}}{\ell_{N,t_0}}. \quad (13)$$
>
> When $s_N/s_C < 1$, as it is here, cognitive employment falls by less than N rises, and therefore by less than Y rises; *the loss in cognitive employment is smaller than the gain in GDP.*

Equation (13) pins down target employment in cognitive and all other occupations once one knows the impact on output and wages. Earlier, Section 2.1.4 provided them only to first

[3] Substituting the first line of Equation (11) into the second, AI raises the average wage if and only if

$$\varepsilon \;>\; \varepsilon^*_t \;\approx\; \frac{1}{s_{L,t_0}}\Big[\, s_{K,t_0} - \sigma \;+\; \psi_t\Big(\frac{1-\rho}{a_t} - (1-\sigma)\Big)\Big], \qquad (12)$$

and labor keeps the fraction $(\varepsilon-\varepsilon^*_t)/(\varepsilon+\sigma/s_{L,t_0})$ of the wage gain a pegged rental rate would give it.

order to develop intuition for how the model works. The following proposition provides the exact solution.

Proposition 1 (The model in closed form)**.** The exact solution for the full model as described so far in closed form is

$$s_{L,t} = 1 - \Big[s_{K,t_0} + s_{L,t_0} \sum_i \psi_{i,t}\, m_i d_{i,t}\, \big(e^{-(1-\sigma)\, a_{i,t}} - \rho_i\big)\Big]\, e^{(1-\sigma)\,\Delta \ln r_t}, \tag{14}$$

$$\tilde{\ell}_{N,t} = -\ln\Big(1 - \sum_i m_i d_{i,t}\, \big[\, 1 - \rho_i\, \psi_{i,t} - (1-\psi_{i,t})\, e^{-(1-\sigma)\, a_{i,t}}\, \big]\Big), \tag{15}$$

$$\Delta \ln w_t = \frac{\Delta \ln s_{L,t} + \tilde{\ell}_{N,t}}{1-\sigma}, \qquad \Delta \ln(Y_t/L_t) \;=\; \Delta \ln w_t - \Delta \ln s_{L,t}, \tag{16}$$

$$\Delta \ln K_t = \ln \frac{1 - s_{L,t}}{s_{K,t_0}} + \Delta \ln(Y_t/L_t) - \Delta \ln r_t, \tag{17}$$

where $\Delta \ln s_{L,t} = \ln(s_{L,t}/s_{L,t_0})$, and the two targets follow from $\tilde{\ell}_{N,t}$ by Equation (13). The rental-rate gap is the unique root of the market for capital,

$$\varepsilon\, \Delta \ln r_t \;=\; \Delta \ln K_t, \tag{18}$$

whose right side, Equation (17), is strictly decreasing in $\Delta \ln r_t$; at $\varepsilon = \infty$ the root is zero. With uniform parameters the two sums are $\psi_t m_t d_t\, [e^{-(1-\sigma)a_t} - \rho]$ and $m_t d_t\, [1 - \rho\psi_t - (1-\psi_t)e^{-(1-\sigma)a_t}]$.

Intuition. Substituting the first-order relations (11) into Equation (13), or expanding Equation (15), gives the two employment group targets in a form in which they are easier to understand:

$$\tilde{\ell}_{N,t} \;\approx\; \underbrace{(1-\rho)\, \psi_t\, m_t d_t}_{\text{net displacement}} \;+\; \underbrace{(1-\sigma)\, (1-\psi_t)\, m_t d_t\, a_t}_{\text{labor released by the gains}}, \qquad \tilde{\ell}_{C,t} \;\approx\; -\frac{s_{N,t_0}}{s_{C,t_0}}\, \tilde{\ell}_{N,t}, \tag{19}$$

where $\tilde{\ell}_{C,t} \equiv \ln(\ell^*_{C,t}/\ell_{C,t_0})$ is the required change in cognitive log employment.

The two terms that determine target employment in the "other" occupations are the same forces that push workers out of cognitive occupations. The first is net displacement: the tasks that automation moves to capital, net of the tasks it creates. The second is $(1-\sigma)$ times the gain on the augmented instances. Each augmented instance needs a_t log points less labor. But because tasks are gross complements, demand for those instances rises by σa_t because they are cheaper, so only the fraction $1-\sigma$ of the labor it saves is released. At $\sigma = 0.5$, half of every gain is spent on more of the same task.

2.2 Innovation

2.2.1 Ideas production and research funding

So far we have held the ideas stock A_t at its no-AI path. This section makes it semi-endogenous, with AI entering through the resources the economy devotes to research. Ideas are produced from research input R_t, in the tradition of Romer (1990) and Jones (1995):

$$\dot{A}_t = \nu R_t^{\lambda} A_t^{\phi_R}, \qquad g_{A,t} = \nu R_t^{\lambda} A_t^{\phi_R - 1}. \tag{20}$$

We set $\lambda \leq 1$ and $\phi_R < 1$: proportional gains become harder to find as the stock of ideas grows.

Research. Research uses final output, as in the lab-equipment model of Rivera-Batiz and Romer (1991). The economy devotes a fraction $\iota_{R,t}$ of GDP to research, funded by a lump-sum tax on households; the government pays the researchers' salaries and their equipment and compute out of that budget.[4] Output is divided as

$$Y_t = C_t + I_t + R_t, \qquad R_t = \iota_{R,t} Y_t, \tag{21}$$

where R_t is research, I_t is the investment, and C_t is consumption.[5]

Impact effects. AI raises research input because the budget is a share of GDP and AI raises GDP: the level channel of Section 2.1 makes the economy more productive, and a given share of a bigger economy is a bigger research budget. Log-differencing $R_t = \iota_{R,t} Y_t$ against the path without AI gives

$$\Delta \ln R_t = \Delta \ln Y_t, \tag{22}$$

where $\Delta \ln Y_t$ is the gap in GDP and the research share, common to the two paths, has dropped out. AI automates research by automating the goods that it uses to produce ideas.

Now consider the impact of AI on growth. Writing $\Delta g_t \equiv g_{A,t} - g$ for the gap in the growth rate and taking log differences of (20),

$$\Delta \ln g_t = \lambda \Delta \ln R_t - (1 - \phi_R) \Delta \ln A_t. \tag{23}$$

This is the key equation of the section. The percent gap in the growth rate is smaller than the percent gap in research inputs, which in turn is the GDP gap of Section 2.1.

> The GDP gap is a gain in the *level* of output, whereas here it applies to a *growth rate* of around one percent per year. A GDP gap of 5 percent therefore raises the growth rate of ideas by less than 5 percent, e.g. raising measured TFP growth from 1.00 to 1.05 percent. The number is even smaller since $\Delta \ln A_t > 0$ and to the extent that $\lambda < 1$. The effects of AI on idea growth here tend to be small.

[4] We allow $\iota_{R,t}$ to rise exogenously over time capturing the rise in R&D as a share of GDP in the U.S. and other economies; see Appendix C.

[5] Given the way we model capital — including the possibility that some is financed from abroad — there is no need to specify the split between consumption and investment.

Accumulation over time. A one-time increase in research input, however, keeps adding to A_t year after year. Because returns to the ideas stock are decreasing ($\phi_R < 1$), a given level of research input R is consistent with growth at the baseline rate g at only one level of the ideas stock: setting $g_{A,t} = g$ in Equation (20) gives $A^* = (\nu R^\lambda / g)^{1/(1-\phi_R)}$, so that $A^* \propto R^{\lambda/(1-\phi_R)}$ and an uplift in research input, held there, moves the ideas stock in the long run by

$$\Delta \ln A^* \;=\; \gamma\, \Delta \ln R, \qquad \gamma \;\equiv\; \frac{\lambda}{1-\phi_R} \;=\; \frac{1}{2.86} \;\approx\; 0.35, \tag{24}$$

where the value for γ is based on Bloom et al. (2020).

Equation (24) therefore already gives the order of magnitude of this channel. On the extreme path, the uplift $\Delta \ln R_t$ reaches 0.28 in 2030, so the ideas stock is headed for a level 10 percent above its no-AI path (and even less in TFP units). That is the most the channel delivers for a one-off change of that size. The realized values by 2030 are smaller still because the long-run gains take decades to cumulate. On the extreme path, the ideas stock is just 0.6 percent above its no-AI path in 2030.

A simplification in our setup is that there is no feedback from ideas to the gains from AI, a_t. Research raises the productivity of everything people do but does not by itself automate any of it. We instead chose to let a_t and ψ_t be parameters of the scenarios themselves. In that sense, the gains from recursive self-improvement (RSI) can be thought of as being hard-coded into a_t.[6]

2.2.2 Measured TFP

The ideas stock A_t is not what an econometrician would measure as TFP. Measured TFP is the Solow residual, $d \ln \mathrm{TFP}_t \equiv d \ln Y_t - s_{L,t}\, d \ln L_t - s_{K,t}\, d \ln K_t$, and in this model it combines the ideas stock with the level channel of Section 2.1. It has a simple closed form. As in Section 2.1.3, because factor payments exhaust output, the residual equals its dual:

$$g_{\mathrm{TFP},t} \;=\; s_{L,t}\, g_{w,t} \;+\; s_{K,t}\, g_{r,t}. \tag{25}$$

Both the AI level term $m_t d_t a_t$ and the growth term $\Delta \ln A_t$ enter the residual with the labor share, its Domar weight, as in Equation (9). The ideas stock lowers the cost of every labor instance by $d \ln A_t$, and the level channel lowers the cost of the affected ones by a_t. To first order, Equation (25) therefore gives the cumulative gap against the no-AI path as

$$\Delta \ln \mathrm{TFP}_t \;\approx\; s_{L,t_0}\left(\Delta \ln A_t \;+\; m_t\, d_t\, a_t\right). \tag{26}$$

[6] This does not have to be the case of course, and this channel merits further exploration in the future.

2.3 Unemployment

We now describe the flows that move employment toward the targets of Section 2.1.5 and study the implications for unemployment, at a monthly time horizon. In normal times a pool of $\bar{U}$ unemployed workers is between jobs at any moment, so that base period employment is $\ell_{C,t_0}+\ell_{N,t_0} = L-\bar{U}$. The targets of Equation (13) are adjusted for this steady-state unemployment, with $\ell_{N,t_0} = (s_{N,t_0}/s_{L,t_0})(L - \bar{U})$, and also add up to $L - \bar{U}$.[7]

2.3.1 Separations and job openings

Workers in group $o \in \{C, N\}$ experience two types of separations: quits, and layoffs caused by displacement. They quit at the rate

$$q_{o,t} \;=\; q_o^X \;+\; q_o^T \,\frac{f_{o,t-1}}{\bar{f}_o} \tag{27}$$

where q_o^X is an exogenous base rate and the second term makes quits rise with job prospects, captured by the group's job-finding rate $f_{o,t-1}$ in the previous month relative to its steady-state value $\bar{f}_o$. In normal times, the quit rate is thus $q_{o,t} = q_o^X + q_o^T \equiv \bar{q}_o$.

Each month, employers in group o compare their current employment with next month's target. On our scenario paths, the cognitive employment target only falls and the other only rises, so the cognitive occupations carry an *employment overhang* and the others an *employment shortfall*, both measured as log gaps,

$$G_{C,t} \;=\; \max\left\{0,\; \ln \ell_{C,t} - \ln \ell^*_{C,t+1}\right\}, \qquad B_{N,t} \;=\; \max\left\{0,\; \ln \ell^*_{N,t+1} - \ln \ell_{N,t}\right\}. \tag{28}$$

The two sides of the market adjust differently. The expanding occupations post openings for a fraction θ^H of any shortfall per month. The shrinking occupations lay off the workers they do not demand at a slowly-adjusting wage, as discussed next. Unemployment is then the outcome of two frictions, a wage that falls too slowly to keep the displaced employed and a search process that takes time.

The cognitive wage and layoffs. Let $N_{C,t}$ be the cognitive labor force attached to the group, the employed plus the cognitive-origin unemployed above their normal pool,

$$N_{C,t} \;=\; \ell_{C,t} + \max\{0,\; U_{C,t} - \bar{U}_C\}, \tag{29}$$

and let $w^c_{C,t}$ be the wage that would clear the cognitive group at $N_{C,t}$, i.e. if no extra C workers were pushed into unemployment. In normal times it is the common wage w_t, but while workers are displaced, the wage required to fully employ them is even lower.

[7] As in Section 2.1.1, every head count in this section (employment, the pool, the attached force, openings, hires, layoffs, and the targets) is a share of the labor force L_t, which grows at n with or without AI, so L without a subscript is the constant labor force in these units and the equations hold as written in shares; rates (q, f, π, θ^H) are per month.

We assume the cognitive wage is slow to adjust, so the cognitive labor market does not clear and there is unemployment. Specifically, the actual cognitive wage paid, $w_{C,t}$, closes only a fraction of its gap to $w^c_{C,t}$ each month, in the manner of Blanchard and Galí (2007):

$$\frac{w_{C,t}}{w_t} = \left(\frac{w_{C,t-1}}{w_{t-1}}\right)^{\xi_m} \left(\frac{w^c_{C,t}}{w_t}\right)^{1-\xi_m}, \qquad \xi_m = \xi^{1/12}, \tag{30}$$

where $\xi \in [0, 1)$ is the rigidity of the cognitive wage per year. At $\xi = 0$ the wage clears the group every month, and as $\xi \to 1$ it is the common wage. This rigidity occurs relative to the common wage w_t; what is rigid is the cognitive discount.

At this wage, the cognitive labor market does not clear, and firms adjust employment to be on their labor demand curve. Because of the search/matching structure of the model, employment is predetermined within the month, so firms reach their demand through separations. Let $\ell^d_{C,t}$ denote the labor firms demand at $w_{C,t}$, and let $E_t = \max\{0, \ell_{C,t} - \ell^d_{C,t}\}$ be the excess employment they carry into the month. Quits from surplus positions are not replaced, and displacements $D_{C,t}$ ("layoffs") remove the rest:

$$D_{C,t} = \max\{0, \; E_t - q_{C,t}\,\ell_{C,t}\}, \tag{31}$$

so that cognitive employment is on the demand curve from the next month on.

The layoffs are job rationing in the sense of Michaillat (2012): while the wage is above its clearing level, firms demand fewer cognitive workers than are attached to the group, and the difference is laid off rather than employed at a lower wage. While rationing binds, the employed cognitive workers are paid their marginal product.

Job openings. Firms post openings when they need to replace quits and to close a fraction θ^H of their employment shortfall. For the cognitive group, the shortfall is demand beyond employment, $Z_t = \max\{0, \ell^d_{C,t} - \ell_{C,t}\}$, typically zero; for the all-other group, it is the log gap $B_{N,t}$ to its target. Job openings $v_{o,t}$ are then

$$v_{C,t} = \frac{\max\{0, \; q_{C,t}\,\ell_{C,t} - E_t\} + \theta^H Z_t}{\bar{\pi}_C}, \qquad v_{N,t} = \frac{\left(q_{N,t} + \theta^H B_{N,t}\right)\ell_{N,t}}{\bar{\pi}_N}, \tag{32}$$

where $\bar{\pi}_o < 1$ is the fraction of postings that employers expect to fill within the month, their normal-times filling rate, and $v_{o,t}$ is a flow of postings per month. At $\theta^H = 1$, the all-other group posts its entire shortfall at once. A θ^H below one spreads the postings over time and stands for the recruiting and training capacity of the expanding occupations.

2.3.2 Matching

Separated workers enter the unemployment pool, and we track the unemployed by group of origin, $U_{C,t}$ and $U_{N,t}$. They search in both groups, but a worker searching outside the group of origin is at a disadvantage, reflecting occupation-specific human capital (Kambourov and

Manovskii, 2009), which is otherwise outside the model: a software engineer does not become an electrician or a nurse at once. We capture this with a discount $\mu \in (0,1]$ on search across groups, so that the *effective search effort* directed at each group is

$$S_{C,t} \;=\; U_{C,t} + \mu\, U_{N,t}, \qquad S_{N,t} \;=\; \mu\, U_{C,t} + U_{N,t}. \tag{33}$$

Hires in group j are given by a matching function combining search effort and openings, of the form proposed by den Haan et al. (2000),

$$H_{j,t} \;=\; \chi\, \frac{S_{j,t}\, v_{j,t}}{\left(S_{j,t}^{\iota} + v_{j,t}^{\iota}\right)^{1/\iota}}, \qquad j \in \{C, N\}, \tag{34}$$

with the standard curvature $\iota = 1.27$ and matching efficiency $\chi \le 1$.[8] The function has constant returns to scale, so the filling rate and hires per unit of effective search depend only on the ratio of openings to searchers, the tightness of the group's market, $\theta_{j,t} \equiv v_{j,t}/S_{j,t}$. The filling rate is $\pi_{j,t} \equiv H_{j,t}/v_{j,t} = \chi\left[1 + \theta_{j,t}^{\iota}\right]^{-1/\iota}$ and hires per unit of effective search are $H_{j,t}/S_{j,t} = \theta_{j,t}\,\pi_{j,t}$. Both lie between zero and χ whatever the tightness, so hires never exceed the searchers directed at a group or its openings. This is why we use this form rather than the Cobb–Douglas function of most of the studies that Petrongolo and Pissarides (2001) survey, which has to be capped to keep hires within these bounds, and the cap puts kinks into the simulated scenario paths. As $\iota \to \infty$ the function becomes $\chi \min\{S_{j,t}, v_{j,t}\}$.

A group hires at the rate $H_{j,t}/S_{j,t}$ per unit of effective search directed at it, and a worker supplies one unit of search in the group of origin and μ units in the other, so that the job finding rates by origin are

$$f_{C,t} \;=\; \frac{H_{C,t}}{S_{C,t}} + \mu\,\frac{H_{N,t}}{S_{N,t}}, \qquad f_{N,t} \;=\; \mu\,\frac{H_{C,t}}{S_{C,t}} + \frac{H_{N,t}}{S_{N,t}}. \tag{35}$$

A laid-off cognitive worker's chance of finding a job is therefore high when all other occupations are posting heavily and μ is close to one, and low otherwise. Employment then evolves through the two kinds of separation (quits and layoffs) and through hires. Each pool of unemployed workers evolves through the same separations less the workers who find jobs:

$$\ell_{o,t+1} \;=\; (1 - q_{o,t})\,\ell_{o,t} \;-\; D_{o,t} \;+\; H_{o,t}, \qquad o \in \{C, N\}, \tag{36}$$

$$U_{o,t+1} \;=\; U_{o,t} \;+\; q_{o,t}\,\ell_{o,t} + D_{o,t} \;-\; f_{o,t}\,U_{o,t}, \qquad o \in \{C, N\}, \tag{37}$$

where $D_{N,t} = 0$, and a worker hired into group j counts as a j worker from then on. Notice that $f_{C,t}U_{C,t} + f_{N,t}U_{N,t} = H_{C,t} + H_{N,t}$ by Equations (33) and (35): every hire leaves the pool and every separation enters it, so $\ell_{C,t} + \ell_{N,t} + U_{C,t} + U_{N,t} = L$ at every date.

New entrants to the labor force join the two groups and the unemployed pool in proportion

[8] Den Haan et al. set $\chi = 1$; we keep χ so that the level of the filling rate can be calibrated separately from the curvature. The elasticity of hires with respect to search effort, $(v_{j,t}/S_{j,t})^{\iota}/\left[1 + (v_{j,t}/S_{j,t})^{\iota}\right]$, is not a parameter of this function but rises with tightness; where openings are scarce, additional searchers add little to hires.

to their sizes, so in shares of L_t the laws of motion (36) and (37) hold. All separations pass through the unemployment pool and so do workers who change groups. Historically, gross occupational flows have dwarfed net flows with much of the adjustment mediated by entry to and exit from the labor force (Pilossoph, 2012; Carrillo-Tudela and Visschers, 2023b; Autor et al., 2014; Cortes et al., 2017). Our focus is on scenarios where displaced workers must cross through the unemployed pool.

Steady state. The block starts from a steady state in which, group by group, hiring replaces quits and each origin's inflow equals its outflow,

$$\bar{H}_o \;=\; \bar{q}_o\,\ell_{o,t_0} \;=\; \bar{f}_o\,\bar{U}_o, \qquad \bar{f}_o \;=\; \sum_j \bar{\mu}_{oj}\,\frac{\bar{H}_j}{\bar{S}_j}, \qquad \bar{\pi}_o \;=\; \chi\Big[1-\Big(\frac{\bar{H}_o}{\chi\,\bar{S}_o}\Big)^{\iota}\Big]^{1/\iota}, \tag{38}$$

for $o \in \{C, N\}$, where $\bar{\mu}_{oo} = 1$ and $\bar{\mu}_{CN} = \bar{\mu}_{NC} = \bar{\mu}$, the search discount of normal times, $\bar{S}_j$ is Equation (33) at the steady-state stocks with $\bar{\mu}$ in place of μ, and $\bar{U}_C + \bar{U}_N = \bar{U}$. The first two conditions split the pool between the two origins, and with it the aggregate finding rate $\bar{f} = \bar{q}\,(L - \bar{U})/\bar{U}$ that the pool implies. The third inverts the matching function at the steady state to give the normal filling rates $\bar{\pi}_o$ that the posting rule (32) uses, and it has a solution as long as normal hiring per unit of effective search is below the ceiling χ.

We set χ so that the employment-weighted mean of $\bar{\pi}_C$ and $\bar{\pi}_N$ is 0.65 per month.[9] The gap between χ and $\bar{\pi}_o$ is the room by which a group's filling rate can rise when many displaced workers search in it at once.

In our normal times the pool is $\bar{U} = 3.8$ percent of the labor force and the quit rates are $\bar{q}_C = 0.63$ and $\bar{q}_N = 1.40$ percent a month, the 0.92 of the economy as a whole split in the proportions of the two groups' separation rates in the Current Population Survey (CPS). The aggregate finding rate these imply is $0.92 \times 96.2/3.8 \approx 0.23$ a month, against 0.22 in the data. The pool splits into 1.76 percent of the labor force of cognitive origin and 2.08 of other origin, that is, 2.9 percent of the one group and 5.4 percent of the other. The filling rates are 0.66 and 0.64, and χ is 0.76.

The search discount behind these numbers is $\bar{\mu} = 0.17$, the value at which the share of job-finders in this steady state who change groups equals one in seven, calibrated to this value in the CPS. The steady state is evaluated at $\bar{\mu}$ in every scenario, so these objects are common to the scenarios and $\bar{U}$ is unaffected by AI by construction.

Two key search parameters capturing how easily cognitive workers can switch occupations and how quickly firms post new vacancies are scenario objects. The modest, substantial, and extreme scenarios set $\mu = 0.17$, 0.08, and 0.04 in Equations (33) and (35) and $\theta^H = 0.10$, 0.25, and 0.50. The modest scenario keeps the normal-times discount and the other two put the displaced cognitive workers at a steeper disadvantage when aiming to switch occupations.

[9] We calibrate it with the daily model of Davis et al. (2013). Applied to the JOLTS hires and openings of 2010–19, that model gives a filling rate of 4.0 percent per working day, a mean vacancy duration of 26 working days, and a share filled within the month of 0.65.

2.3.3 Summary of the model with unemployment

Our simulated economy features unemployment that initially rises as the labor market adjusts to AI. Workers in the cognitive group are paid the sticky wage $w_{C,t}$, while workers in all other occupations are paid their marginal product $w_{N,t}$. Table A.1 states the full system of equations that the simulations solve to characterize the economy with AI.

3. Calibration

We start by summarizing the key parameters in our model and the role that they play:

- The product of the extent of capabilities, m_t, and the diffusion parameter, d_t, is the share of task instances performed with AI. These parameters always enter as a product.
- The magnitude through which AI affects task productivity is a_t. The product $m_t d_t a_t$ determines the increase in TFP and therefore the overall rise in GDP that results from AI.
- The parameter ψ_t is the fraction of AI-performed instances that are automated and shift to capital, while $1 - \psi_t$ are augmented and remain with cognitive workers.
- The reinstatement fraction ρ measures the rate at which new tasks arrive and shift work back from capital to cognitive workers. This is another way in which AI may augment cognitive labor.
- The elasticity of capital supply governs how the productivity gains are split between capital and labor.
- The rigidity of the cognitive wage governs whether the cost to cognitive workers takes the form of lower wages or of unemployment.

The labor market in normal times — how large the two groups are, how large the pool is, how the two groups' separation rates compare, and how often a job-finder crosses from one group to the other — is measured from the Current Population Survey. The level of quits is pinned down by the amount of unemployment and the job finding rate so that the steady state matches the U.S. labor market in 2025. The AI objects start from measurements of AI use and adoption in mid-2026, and their values in 2030 are scenario assumptions.

Our paper considers three main scenarios, ordered by the overall extent to which AI transforms the economy by 2030: modest change, substantial change, and extreme change. The three scenarios differ along two key dimensions. The first is how quickly AI capability and its use advance, captured by the 2030 values of m_t, d_t, and a_t. The second is how disruptive a given amount of AI is for workers, captured by the automation share, the reinstatement ratio, the search discount, and the posting speed. Table 1 lists every number with its source. The three scenarios are illustrative. Our framework allows one to set each scenario object and trace the paths that follow. Section 3.5 does this for the representative answers of a survey of US adults.

Table 1: Calibration: parameters and scenario values.

Symbol	Meaning	Modest	Substantial	Extreme	Source or basis
A. Technology, factor markets, and ideas production (Sections 2.1.1 and 2.2); common to the scenarios					
σ	elasticity of substitution across tasks		0.5		gross complements (Acemoglu and Restrepo, 2022; Humlum, 2019; Jones and Tonetti, 2026)
s_{L,t_0}	base-period labor share		0.60		conventional value for the U.S.
$s_{C,t_0}/s_{L,t_0}$	cognitive share of employment		0.624		CPS 2025 (U.S. Bureau of Labor Statistics, 2026b): share of employees in SOC major groups 11–29, 41, 43
ε	elasticity of capital supply		3		half the long-run wealth elasticity of Moll et al. (2022)
$\bar{r}$; δ	no-AI gross rental rate; depreciation, per year		0.115; 0.05		Moll et al. (2022); net return 6.5 percent, $K/Y = 3.5$
t_0	base period		2024		
	anchor date of the paths		mid-2026		
	date of the scenarios' 2030 values		2030.0		
h	period length		1 month		
λ	returns to research input		1		no duplication
$1-\phi$	fishing out, TFP units, research input in researcher-equivalents		3.1		Bloom et al. (2020)
$1-\phi_R$	fishing out, labor-augmenting units, research input in goods		2.86		$s_{L,t_0}(1-\phi)+\lambda$, Equation (40)
g; g_A	no-AI growth of the ideas stock; of measured TFP, $s_{L,t_0}\, g$, per year		0.0167; 0.010		postwar TFP growth
n	growth of the labor force, per year		0.0033		so that GDP grows at 2 percent without AI; BLS projections put labor-force growth near 0.4
ι_{R,t_0}; g_ι	research share of GDP in 2024; its growth, per year		0.035; 0.028		R&D share in the national accounts; Equation (41) in Appendix C, so that A_t grows at g without AI ($\iota_{R,2030} = 0.041$)
B. AI in the production of goods and services (Section 2.1.2)					
m_t	affected mass, mid-2026 anchor		0.14		observed exposure, averaged over occupations (Massenkoff and McCrory, 2026)
m_{2030}	affected mass, 2030	0.2	0.3	0.5	scenario assumptions: the low end, middle, and near the high end of rated feasibility (Eloundou et al., 2024)
$\bar{m}$;	ceiling on affected mass		0.624		all cognitive work

continued on the next page

Table 1, continued

Symbol	Meaning	Modest	Substantial	Extreme	Source or basis
d_t	diffusion share, mid-2026 anchor		0.10		firm AI use, Business Trends and Outlook Survey and its 2026 AI supplement (U.S. Census Bureau, 2026; Bonney et al., 2026)
d_{2030}	diffusion share, 2030	0.2	0.4	0.6	scenario assumptions
$\bar{d}$	ceiling on diffusion		1		every instance
a_t	log gain per instance, mid-2026 anchor	0.30	0.35	0.45	scenario assumptions: all-in gains in trials and at the frontier (Brynjolfsson et al., 2025; Huang et al., 2025; Demirer et al., 2026)
g_a	slope, per year	0	0.028	0.10	scenario assumptions which imply $a_{2030} = 0.30$ / 0.45 / 0.80
C. Disruptiveness: automation, new tasks, and search once AI arrives (Sections 2.1.2 and 2.3.2)					
ψ_t	automation share, held constant	0.50	0.75	0.90	scenario assumptions: on the low end similar to current automation-like share of use (Appel et al., 2025)
ρ	reinstatement ratio	0.50	0.25	0	scenario assumptions: Acemoglu and Restrepo (2019) estimated $\rho = 0.5$
μ	search discount on the path	0.17	0.08	0.04	scenario assumptions: $\bar{\mu} = 0.17$ is the value consistent with normal-times occupational switching estimated from CPS
θ^H	posting speed, per month	0.10	0.25	0.50	scenario assumptions
ξ	rigidity of the cognitive wage, per year		0.50		wage-setting evidence (Section 3.4)

D. The labor market in normal times (Section 2.3); common to the scenarios			
$\bar{q}$	normal quit rate, per year	0.11	implied by the pool and the CPS finding rate, rounded (Section 3.2)
$q_o^T/\bar{q}_o$	share of quits that responds to job prospects	0.55	elasticity of JOLTS quits to the CPS finding rate (U.S. Bureau of Labor Statistics, 2026a) (Section 3.2)
$\bar{q}_C/\bar{q}$; $\bar{q}_N/\bar{q}$	relative separation rates by group	0.69; 1.52	IPUMS-CPS 2010–19 (Flood et al., 2025)
$\bar{U}$	normal search pool, share of L	0.038	CPS 2025 (U.S. Bureau of Labor Statistics, 2026b)
$\bar{\mu}$	search discount in normal times	0.17	CPS 2010–19 switching matrix from the replication files of Carrillo-Tudela and Visschers (2023b), corrected by their method (Section 3.2)
ι	matching curvature	1.27	den Haan et al. (2000)
$\bar{\pi}_o$, mean	filling rate, per month	0.65	JOLTS 2010–19, computed using the method of Davis et al. (2013)

Note: A value spanning the three scenario columns is common to the scenarios. The ceilings of the logistic paths are their natural maxima ($\bar{m} = s_{C,t_0}/s_{L,t_0}$, $\bar{d} = 1$), and slope and midpoint follow from 2026 anchor, 2030 value, and ceiling by Equation (8′).

3.1 Parameters related to production functions

Tasks are gross complements with $\sigma = 0.5$, and labor earns 60 percent of income before AI.

The cognitive occupations are the twelve SOC major groups that cover management, professional, sales, and office work, and all other occupations are the remaining ten.[10] The cognitive occupations employ 62.4 percent of workers in the 2025 CPS, and they are where AI is used at work: the observed-exposure measure described next is concentrated almost entirely in them (Section 3.3). Of course, any two-way split is a simplification, since exposure varies widely and many occupations involve some cognitive work.

The observed-exposure measure of Massenkoff and McCrory (2026) scores each occupation by the share of its task time spent on tasks with significant work-related Claude use, with assistive use counted at one half. Claude is one of several AI assistants in use at work, so the measure is a proxy for AI use generally rather than a census of it. Its concentration in the cognitive occupations matches survey-based measures of generative AI use (Bick, Blandin, and Deming, 2024) and the task-feasibility ratings of Eloundou et al. (2024). Recall from Section 2.1.2 that the

[10] By 2018 SOC major group, the cognitive occupations are Management (11); Business and Financial Operations (13); Computer and Mathematical (15); Architecture and Engineering (17); Life, Physical, and Social Science (19); Community and Social Service (21); Legal (23); Educational Instruction and Library (25); Arts, Design, Entertainment, Sports, and Media (27); Healthcare Practitioners and Technical (29); Sales and Related (41); and Office and Administrative Support (43). All other occupations are Healthcare Support (31); Protective Service (33); Food Preparation and Serving Related (35); Building and Grounds Cleaning and Maintenance (37); Personal Care and Service (39); Farming, Fishing, and Forestry (45); Construction and Extraction (47); Installation, Maintenance, and Repair (49); Production (51); and Transportation and Material Moving (53). The scenario explorer labels the two groups “knowledge workers” and “all other workers.”

affected tasks must fit inside this group, so the ceiling of the affected mass is $\bar{m} = s_{C,t_0}/s_{L,t_0}$ in every scenario: in the long run all cognitive work is within AI's reach, and the scenarios differ in how much of it is within reach by 2030.[11]

The ideas parameters are those of Section 2.2: no duplication, fishing out of 3.1 from Bloom et al. (2020) in TFP units and researcher-equivalents, which is 2.86 in the model's labor-augmenting units with research input in goods, growth of the ideas stock of 1.67 percent a year without AI and hence TFP growth of one percent, and a research budget of 3.5 percent of GDP in 2024 whose share rises at 2.8 percent a year, with or without AI (Appendix C).

We set the elasticity of capital supply at $\varepsilon = 3$, which is higher than five years of saving at historical rates would imply. We lean toward a relatively high value because the capital that performs automated cognitive work is disproportionately compute, financed in a world market and built in a year or two.

3.2 Normal times

Panel D is the labor market before AI. Nearly every number comes from the Current Population Survey (CPS).

The 2025 CPS annual averages (U.S. Bureau of Labor Statistics, 2026b) give the two groups' shares of employment and the unemployed pool. The IPUMS-CPS matched monthly files of 2010–19 (Flood et al., 2025) give the flows. Each month, on average, 22 percent of unemployed workers found a job, and 0.84 percent of employed cognitive workers and 1.84 percent of workers employed in all other occupations became unemployed. The calibration uses the ratio of these two rates. The same files record where job-finders go. We tabulate these moves from the replication files of Carrillo-Tudela and Visschers (2023a,b) and remove the occupational switches that are coding errors by their method: 19 percent of job-finders from cognitive occupations took a job in the all-other group, and 11 percent of job-finders from all other occupations took a cognitive job. Overall, 14.3 percent changed group.

Three parameters follow from these facts. First, the quit rate. In normal times hires equal quits, and hires are the job-finding rate times the unemployed pool, so the CPS finding rate and the pool imply a quit rate of $0.219 \times 3.84/96.16 = 0.875$ percent of employment a month, or 0.105 a year. We round it to 0.11 a year, or 0.92 percent a month, and at that value the finding rate in the model's steady state is 0.23 a month. Splitting the quit rate into the ratio of the two separation rates gives 0.63 percent a month for cognitive workers and 1.40 for all other workers.

Second, the search discount $\bar{\mu}$. By Equation (35), a cognitive job-finder's odds of landing in the all-other group are $\bar{\mu}$ times that group's hiring rate relative to the cognitive group's, and the odds that a job-finder from the all-other group lands in a cognitive job are $\bar{\mu}$ times the inverse of that ratio. Multiplying the two odds cancels the ratio and leaves $\bar{\mu}^2$. The occupational switching

[11] The slope that takes the affected mass from its mid-2026 estimate m_{2026} to its 2030 value m_{2030}, three and a half years later, as in equation (8), is

$$\kappa_m = \frac{1}{3.5} \ln\left[\frac{\bar{m} - m_{2026}}{m_{2026}} \cdot \frac{m_{2030}}{\bar{m} - m_{2030}}\right], \tag{8'}$$

and likewise for κ_d.

matrix gives odds of $19/81 = 0.23$ and $11/89 = 0.12$, so $\bar{\mu} = \sqrt{0.23 \times 0.12} = 0.17$. We apply the same cross-occupation search discount in both directions, so the asymmetry between 19 and 11 percent reflects differences in hiring rates.

Third, the split of quits into a part that responds to job prospects and a part that does not (Equation (27)). Over 2001–19 the JOLTS quits rate moved with the CPS job-finding rate with an elasticity of 0.53 (U.S. Bureau of Labor Statistics, 2026a,c), which we round to 0.55.

Finally, the steady state of Equation (38) determines hires, the share of unemployed workers, and the finding rate for each group of origin. The remaining parameter is the matching efficiency χ, which we set so that the employment-weighted mean of the two groups' filling rates is 0.65, a value we compute from JOLTS 2010–19 (Section 2.3.2).

3.3 The pace of AI

All three scenarios start from the same mid-2026 readings of m_t and d_t and differ in their 2030 values. The slope between the two follows from Equation (8′); because each path is a logistic through the anchor with its own slope, the scenarios coincide at mid-2026. Before, they are within a few tenths of a percent of GDP of one another. The affected mass anchor is the observed-exposure measure of Section 3.1. It averages 0.22 in the cognitive group, 0.01 in the other, and 0.14 overall, so $m = 0.14$. The figure is conservative, since assistive use counts half, thinly used tasks count zero, and the data are from late 2025.

The initial value for diffusion draws on the Census Bureau's Business Trends and Outlook Survey (U.S. Census Bureau, 2026; Bonney et al., 2026). In late 2025, 18 percent of firms used AI (32 percent employment-weighted), and in 23 percent of firms workers used generative AI in their own tasks (Bonney et al., 2026). Firm shares overstate the share of *instances* done with AI, because adopters use it on few tasks and on only some instances of each. We therefore set the initial $d = 0.10$.

For 2030, we assume that the affected task mass is 0.2, 0.3, or 0.5, spanning the range of rated feasibility in Eloundou et al. (2024). Diffusion reaches 0.2, 0.4, or 0.6, below the recent rapid adoption of frontier LLMs (Bick, Blandin, and Deming, 2024) but also potentially reflecting the fact that historically diffusion has been much slower (Griliches, 1957; David, 1990; Comin and Hobijn, 2010).

The productivity gain per instance starts at 0.30–0.45 in logs, reflecting estimates from field trials and frontier firms (Brynjolfsson et al., 2025; Huang et al., 2025), about a quarter of the gain of 1.6 estimated from conversations with Claude (Tamkin and McCrory, 2025). The scenarios also differ in whether that gap closes. In the modest scenario, the productivity gain per task is fixed, whereas in the extreme scenario it reaches 0.80 by 2030.

3.4 Disruptiveness

The parameters in Panel C govern the disruptiveness of AI. The automation share ψ is the fraction of AI-performed instances done by AI outright rather than with a worker. The Anthropic

Economic Index classifies conversations by mode of use (Appel et al., 2025): about half of use on Claude.ai and about three quarters on the API looks like automation (the user delegates and the model executes). So 0.5 describes chat use today, and 0.9 a world in which agentic use is the norm.

The reinstatement ratio ρ is the mass of new labor tasks created per unit of tasks automated. In the decomposition of Acemoglu and Restrepo (2019), new tasks roughly kept pace with automated ones over 1947–87 and ran at about half of them over 1987–2017. The modest scenario's 0.5 is thus similar to recent history. The extreme scenario's 0 assumes that AI does not create any new tasks for humans to perform.

The search discount μ and the posting speed θ^H govern how quickly displaced workers are re-employed. Neither can be estimated precisely, given that AI has not yet displaced a significant share of workers. The re-employment of displaced workers in past episodes is the closest evidence. Using CPS data on employment–unemployment–employment episodes, we estimate the search discount to be approximately $\mu = 0.17$ and use this value in the modest scenario.

However, there are at least two reasons to think that workers displaced by AI will find it hard to find jobs in other occupations than historical data may suggest. First, there is significant heterogeneity in crossing probability within cognitive occupations. For example, unemployed sales and office workers take a job outside the cognitive occupations 28 percent of the time while managers and professionals do so only 10 percent of the time. A search discount fitted to the latter group would be about $\mu = 0.09$. Second, when a large part of a group searches at once, for jobs the other occupations have not yet created, crossing is presumably even harder. The substantial scenario therefore sets $\mu = 0.08$ and the extreme scenario $\mu = 0.04$.

The posting speed, θ^H, is an assumption: it closes a tenth, a quarter, and a half of the all-other group's shortfall per month in the three scenarios, respectively. Both frictions matter for the same reason: unemployment rises above its normal level by the number of workers the cognitive occupations have released less the number all other occupations have absorbed, and μ and θ^H set the pace of absorption.

The rigidity ξ sets how fast the cognitive wage moves toward its clearing level: the gap between the two, both relative to the common wage, shrinks by the factor ξ a year. There is no agreed value of this parameter in the literature, and it may depend on circumstances. For workers who keep their jobs, wages are rigid, especially downward: base wages reset about once a year and nominal cuts are rare (Barattieri et al., 2014; Grigsby et al., 2021), so a real wage that falls only through inflation takes a decade to close a large gap. This suggests ξ near 0.9. On the other hand, for workers who change jobs, wages are flexible: new hires' wages move almost one for one with productivity (Haefke et al., 2013), and displaced workers who are re-employed take losses of ten to twenty percent, concentrated on those who change occupation (Jacobson et al., 1993; Davis and von Wachter, 2011; Huckfeldt, 2022; Braxton and Taska, 2023).

We set $\xi = 0.5$ (a half-life of one year) in all three scenarios, motivated by the following reasons. The shock is large, visible, and permanent, and rigidity gives way in such episodes. For example, nominal cuts became common in 2009 and 2020, and during the inflation of 2021–23,

Figure 1: When will AI be able to do tasks as well as skilled professionals? US adults.

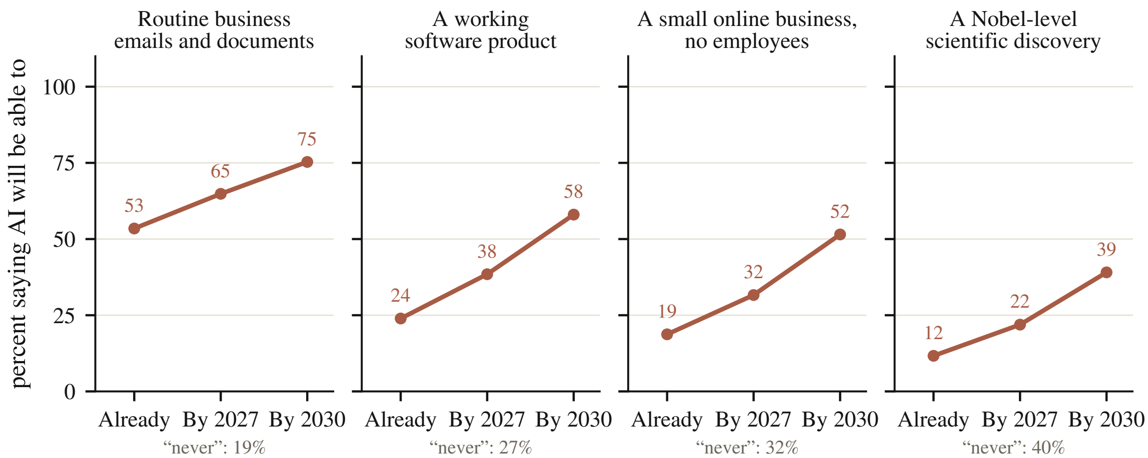

Note: We display four of the eight tasks here, in order of difficulty as respondents rank them. Each line is the cumulative share of respondents from a Morning Consult survey fielded in August 2026 saying AI already can do a task, will by 2027, or will by 2030.

real wages fell by several percent within two years as nominal wages lagged prices. A stickier wage (higher ξ) means more unemployment and a smaller wage decline; we report $\xi = 0.75$ and 0.9 in Section 4.6.

3.5 What do US adults expect?

To compare the three scenarios to people's expectations, we surveyed a representative sample of 10,980 US adults via Morning Consult on August 11–23, 2026. We asked five questions: (a) when AI will be able to do each of eight named tasks (from writing routine business emails to a Nobel-level discovery) as well as a skilled professional, (b) what share of the tasks AI can do it will actually be used for, (c) how much faster an AI-suited task gets done, (d) whether it will do the work alone or alongside a person, and (e) how long a worker displaced by AI will take to find a job in a new occupation. Appendix B describes the survey in further detail.

Figure 1 illustrates the results of our question on AI capabilities. On what AI can already do today, 24 percent of adults say it can build and maintain a software product, and 19 percent say it can run a small online business. Views on how fast AI improves are split. By 2030, 52 percent expect AI to run an online business with no employees and 39 percent expect a Nobel-level discovery, while 40 percent say such discoveries will never happen.

We map each respondent's answers into the model's five main scenario parameters. Capability, m_{2030}, is derived from the share of the eight tasks the respondent expects AI to handle by 2030, placed on a scale from zero to 95 percent of knowledge work (corresponding to all eight tasks) and multiplied by the knowledge-work share of employment. Knowledge work is the scenario explorer's name for the work of the cognitive occupations. Adoption, d_{2030}, and automation, ψ_{2030}, are the shares stated by survey respondents. The productivity gain, a_{2030}, is

Table 2: The five parameters implied by the survey.

		US adults		
Object	Survey question	median	[25th, 75th]	Modest / Subst. / Extreme
m_{2030}	share of eight tasks AI will do as well as a professional by 2030, placed on a scale of knowledge work	0.44	[0.20, 0.59]	0.2 / 0.3 / 0.5
d_{2030}	of the tasks AI can do, share people will actually use it for at work in 2030	0.40	[0.22, 0.61]	0.2 / 0.4 / 0.6
ψ	for five tasks, whether AI will do it alone or with a worker	0.47	[0.25, 0.65]	0.50 / 0.75 / 0.90
a_{2030}	how long an AI-suited task takes with AI versus without, in logs	0.44	[0.09, 1.02]	0.30 / 0.45 / 0.80
μ	months for a worker displaced by AI to find a job in a new occupation (three months = 0.17)	0.064	[0.025, 0.111]	0.17 / 0.08 / 0.04

Note: Answers of survey respondents mapped to model parameters according to the rules described in the text. Quantiles of binned items interpolated within answer bin (Appendix B). The last column repeats Table 1.

set to the log of the reported speed-up. Respondents are told that a job search in normal times takes about three months and are asked how long a worker displaced by AI in 2030 would take to find a job in a new occupation. The search discount scales its normal-times value, $\bar{\mu} = 0.17$, by the ratio of three months to the answer (Appendix B). A response of "not sure" is excluded item by item. Table 2 reports medians and inter-quartile ranges for all five parameters.

We set all other parameters to the substantial change scenario's values. In the associated scenario explorer a user can set these values freely. Further details on the parameter mappings are described in Appendix B.

How do these answers compare with our scenarios? The median respondent's answers are mixed. They expect AI to handle six of the eight tasks by 2030, putting the AI capability parameter to $m_{2030} = 0.44$, between the substantial and the extreme scenario. They expect it to be used on 40 percent of what it can do and to cut the time on an AI-suited task by a third. They also expect half of that use to be automation, the modest scenario's share. And they expect re-employment to take about eight months, slower than the substantial scenario assumes and faster than the extreme.

There is, however, a wide dispersion of views. For example, approximately 30 percent of respondents expect AI to save no time at all on a task suited to it, while 49 percent expect it to cut the time at least in half.

Table 3: The three scenarios in 2030.

	No AI	Modest	Substantial	Extreme
A. Output				
GDP, pct. above the no-AI path	0	1.6	8.3	32.4
GDP, index with 2024 = 100	112.7	114.5	122.1	149.3
GDP growth, pct. per year	2.0	2.4	5.4	15.4
B. Factor prices and income shares				
Average wage, pct. above the no-AI path	0	0.7	2.1	9.7
cognitive occupations, $w_{C,t}$	0	0.4	−0.3	−11.5
all other occupations, $w_{N,t}$	0	1.1	5.9	33.6
Net return to capital $r_t - \delta$, pct. per year	6.5	6.6	7.0	8.3
Capital stock, pct. above the no-AI path	0	2.3	13.8	56.3
Labor share, pct. of income	60.0	59.4	56.1	45.2
Capital share, pct. of income	40.0	40.6	43.9	54.8
Labor income, pct. above the no-AI path	0	0.6	1.4	0.5
wage bill of cognitive occupations	0	−0.3	−4.6	−31.0
Capital income, pct. above the no-AI path	0	3.1	18.9	81.4
C. The labor market				
Cognitive employment, pct. change since mid-2026	0	−0.5	−3.9	−21.5
Unemployment rate, cognitive workers, pct.	2.9	2.9	4.5	17.9
Unemployment rate, all workers, pct.	3.8	3.9	4.6	11.9
D. Technology: the two channels				
Measured TFP, pct. above the no-AI path	0	0.7	3.1	13.4
Measured TFP growth, pct. per year	1.0	1.2	2.3	7.3
Ideas stock A_t, pct. above the no-AI path	0	0.07	0.20	0.61
Growth of the ideas stock, pct. per year	1.67	1.69	1.76	2.02

Note: The three scenarios at the start of 2030, the date convention of Section 2.1.2. Growth rates are log changes over the twelve months to 2030. Labor income is the sum of the two groups' wage bills at the wages and employment shown; capital income is the rental rate times the capital stock.

4. Results

The results from our three scenarios are shown in two different ways. Table 3 shows the key results for 2030. Figures 2 to 4 plot the main economic variables over time since 2025, showing how the effects of AI cumulate. We organize our discussion by considering one scenario at a time.

4.1 The modest change scenario

In the modest change scenario, AI is a small technology. The economy continues on a "normal" path with AI making changes around the margins.

In this scenario in 2030, just 4 percent of the tasks in the economy have been impacted by

Figure 2: GDP and its growth rate in the three scenarios, 2025–2030.

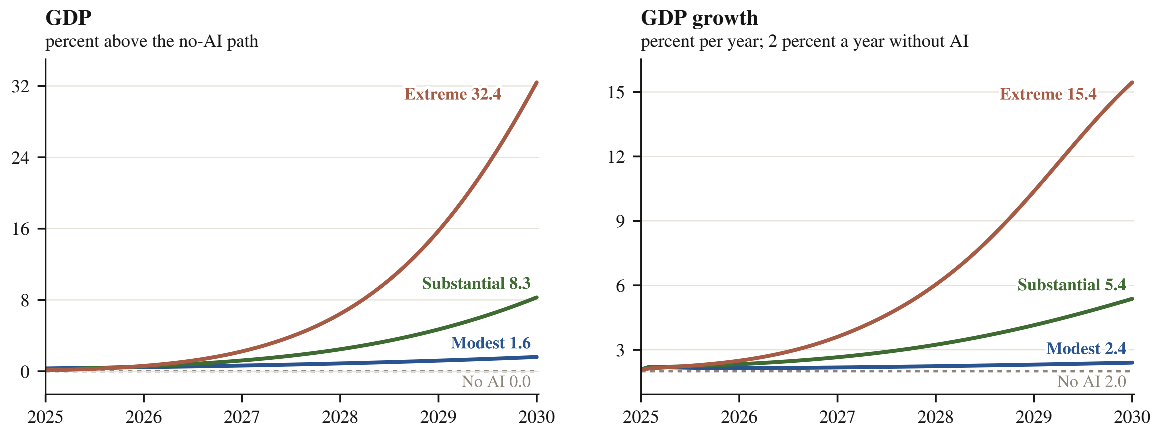

Note: Left: GDP, actual output in percent above the no-AI path. Right: the growth rate of GDP over the preceding twelve months. The dashed line is the economy without AI; each line is labeled with its value in January 2030.

AI, in part because of modest capabilities and in part because of slow diffusion ($m_{2030} = 0.20$ and diffusion is $d_{2030} = 0.20$). AI raises productivity by around 35% on tasks that it impacts ($a_{2030} = 0.30$). Half of the affected instances are automated and half are augmented ($\psi_{2030} = 0.50$). Finally, for every two tasks that AI automates, one new task is introduced to the economy and performed by cognitive labor ($\rho = 0.5$).

GDP is 1.6 percent above its no-AI path in 2030, an extra ten months of ordinary growth accumulated over four years, and the economy is growing at 2.4 percent a year rather than 2 percent. The average wage is 0.7 percent higher, with the cognitive wage up 0.4 percent. The net return to capital is 6.6 percent rather than 6.5, the capital stock is 2.3 percent larger, and the labor share has fallen from 60.0 to 59.4 percent of income. Cognitive employment is 0.5 percent below its mid-2026 level. The unemployment rate is 3.9 percent against a normal level of 3.8: the modest change scenario has only small labor market effects.

4.2 The substantial change scenario

In the substantial change scenario, AI is a pivotal technology, more significant than the internet. It has substantial macroeconomic effects, both on GDP as a whole and on the cognitive occupations.

In this scenario in 2030, 12 percent of the tasks in the economy — roughly 20 percent of the tasks performed by cognitive workers in 2025 — have been impacted by AI ($m_{2030} = 0.30$ and diffusion is $d_{2030} = 0.40$). On tasks that it impacts, AI raises productivity by approximately 57% ($a_{2030} = 0.45$), and three quarters of the affected instances are automated ($\psi_{2030} = 0.75$) while one quarter are augmented. Finally, for every four tasks that AI automates, one new task is introduced to the economy and performed by cognitive labor ($\rho = 0.25$).

By 2030, GDP is 8.3 percent above its no-AI path, and growth over the twelve months to 2030

Figure 3: Factor prices and the labor share in the three scenarios, 2025–2030.

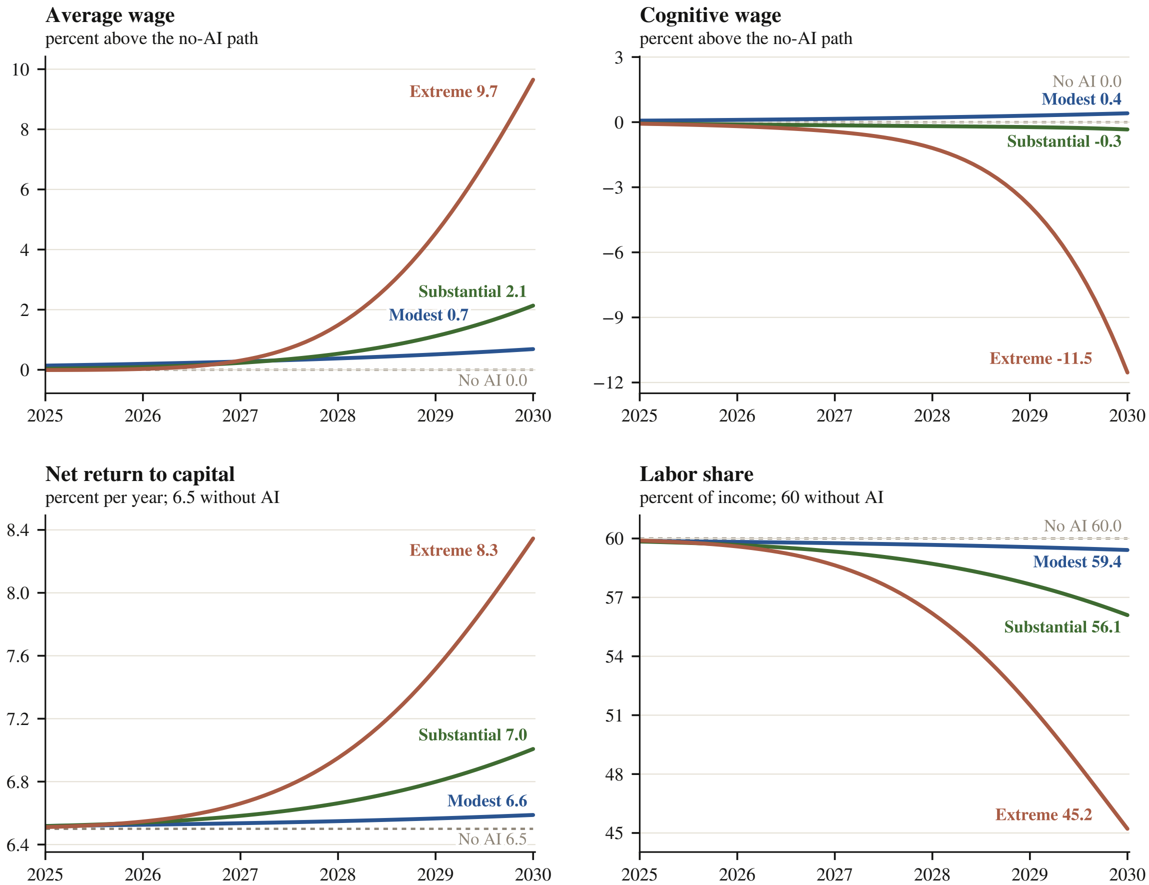

Note: Dashed lines are the economy without AI; each line is labeled with its value in January 2030.

is 5.4 percent a year against 2 percent without AI. For comparison, the fastest GDP growth rate during the 1990s dot-com boom was 4.7 percent in 1999.[12]

How is this income distributed? AI in 2030 has raised the average wage by 2.1 percent, much less than its effect on GDP. Two nuances are behind this statistic. First, the impact on cognitive occupations that are directly displaced differs sharply from the effects on the rest of the economy. Cognitive wages are below the no-AI path by 0.3 percent in 2030, whereas wages in the rest of the economy are higher by 5.9 percent.

Second, because wage growth is slower than GDP growth, factor income has shifted toward capital. In the substantial scenario the labor share falls from 60 percent of income in 2026 to 56 percent in 2030, and the capital share rises correspondingly from 40 to 44 percent: a shift of four points in four years, only slightly less than the entire decline in the U.S. labor share over the four decades after 1980 (Karabarbounis and Neiman, 2014).

[12] https://fred.stlouisfed.org/series/GDPC1.

Figure 4: The labor market in the three scenarios, 2025–2030.

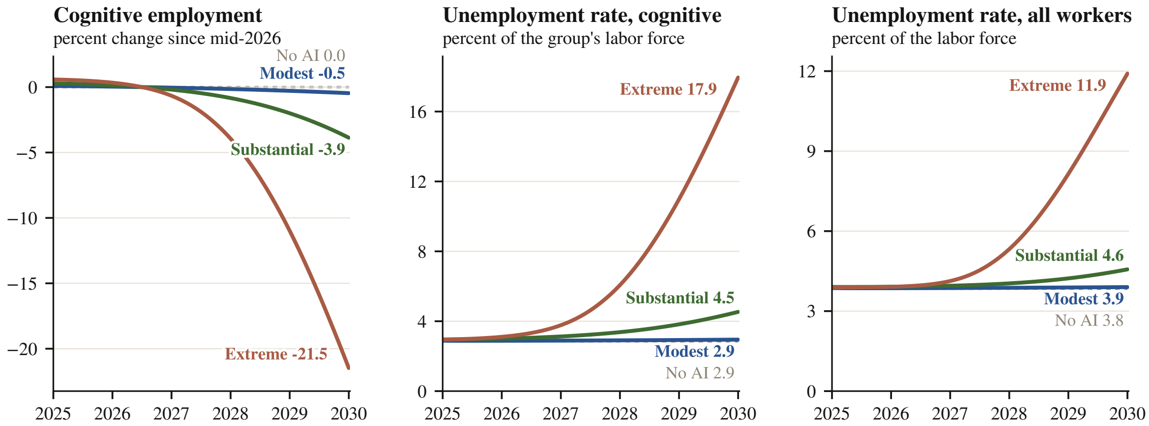

Note: Dashed lines are the economy without AI, in which both rates stay at their normal levels of 2.9 and 3.8 percent; each line is labeled with its value in January 2030.

Finally, there are significant changes related to employment and the reallocation of labor. Cognitive employment is 3.9 percent below its mid-2026 level, and employment in all other occupations has risen by 4.6 percent. Because reallocation takes time, the unemployment rate of cognitive workers rises from 2.9 percent in mid-2026 to 4.5 percent in 2030, more than a 50 percent increase.

4.3 The extreme change scenario

In the extreme scenario, AI transforms the economy. The macroeconomic consequences go well beyond any event in history, both in terms of magnitudes and in terms of the speed at which change happens.

Recall the basic inputs in this scenario: By 2030, 30 percent of the tasks in the economy — roughly half of the tasks performed by cognitive workers in 2025 — have been impacted by AI ($m_{2030} = 0.50$ and diffusion is $d_{2030} = 0.60$). AI is a powerful technology in this scenario, more than doubling productivity on the tasks it affects ($a_{2030} = 0.80$ and $\exp(0.8) \approx 2.2$). Moreover, 90% of the affected instances are automated ($\psi_{2030} = 0.90$) while just 10% are augmented. Finally, the reinstatement effect, which assigned some new tasks to cognitive labor in the other two scenarios, is turned off ($\rho = 0$).

As a result of these changes, by 2030, GDP is 40 percent higher than in mid-2026, putting it 32 percent above the no-AI path. GDP growth in 2030 accelerates to an unprecedented 15.4 percent per year. If this rate were sustained, per capita incomes would double every five years, compared with doubling every 35 years in the US over the past century.

The effects on wages and the distribution of income are equally transformative. The average wage is 9.7 percent above the no-AI path. But this masks the differential impacts by occupation. Cognitive wages are 11.5 percent *below* the no-AI path. And because wage growth in the no-AI scenario is around 2 percent per year, this means that cognitive wages are slightly lower in 2030 than in mid-2026. In contrast, wages in all other occupations, where labor is now comparatively scarce, are 33.6 percent higher.

The factor income distribution also changes dramatically. The automation of nearly half the tasks performed by cognitive labor shifts income to capital and away from labor. The capital share rises from 40% to 55% while the labor share declines from 60% to 45%. A full 15 percent of GDP moves from being paid to labor to being captured as a return on investment. The net return to capital rises from 6.5 to 8.3 percent, an increase of about 28 percent.[13]

By 2030, cognitive employment is 21.5 percent below its mid-2026 level, and the unemployment rate for workers who began in cognitive occupations is a stunning 17.9 percent, far above any postwar rate in the United States. The economy-wide unemployment rate rises to 11.9 percent. By comparison, annual unemployment rates peaked at slightly below 10 percent in the U.S. after the global financial crisis and at about 8 percent during the Covid-19 recession.

Who gains and who loses. In the extreme scenario, GDP is a third larger by 2030, but cognitive workers earn less, and more of them are unemployed. With GDP this much higher, transfers could in principle leave everyone better off. How much redistribution would be required, and what would it need to look like?

Total labor income in 2030 in the extreme scenario is almost exactly what it would have been without AI: the labor share falls by a quarter while GDP rises by a third, and $0.45 \times 1.32 \approx 0.60$. All of the increase in GDP therefore accrues as capital income, which is 81 percent above its no-AI path (panel B of Table 3). Within labor, the wage bill of the cognitive occupations (counting the unemployed at zero) is 31 percent below its previous path as wages are 11.5 percent lower paid on 21.5 percent fewer jobs. The wage bill of all other occupations is higher by nearly the same amount. The economy's gain is thus almost three times the cognitive occupations' loss, so in the textbook sense, the gains could compensate the losses: a transfer of about 9 percent of GDP, roughly the size of Social Security and Medicare combined, would hold cognitive workers' income at its no-AI level and leave the rest of the economy more than 20 percent ahead. Transfers of that scale in response to technological change have no precedent, and the experience of past displacements, such as the regional effects of Chinese import competition (Autor, Dorn, and Hanson, 2013), is that it mostly does not happen on its own. The question of what institutions could lead to shared prosperity in this scenario is beyond our framework.

4.4 The scenarios implied by survey responses

Recall that Table 2 showed the results of the survey we conducted with U.S. adults. Table 4 reports the model outcomes implied by each respondent's answers. The survey asks about AI's capabilities, its use, and the speed of re-employment, and the model works out what those answers imply for output, wages, and unemployment.

[13] The current version of our model possesses an upward-sloping supply of capital but does not feature an Euler equation connecting the return to capital to the growth rate of the economy. This connection could raise the return to capital substantially and will be included in a future version of the model.

Table 4: The economies in 2030 that the model implies from the survey answers.

	US adults				
	median	[25th, 75th]	Modest	Subst.	Extreme
A. Output					
GDP, pct. above the no-AI path	8.6	[3, 19]	1.6	8.3	32.4
GDP growth, pct. per year	5.3	[3, 10]	2.4	5.4	15.4
B. Factor prices and income shares					
Average wage, pct. above the no-AI path	2.6	[0, 13]	0.7	2.1	9.7
cognitive occupations, $w_{C,t}$	0.6	[−2, 7]	0.4	−0.3	−11.5
all other occupations, $w_{N,t}$	6.4	[1, 22]	1.1	5.9	33.6
Net return to capital $r_t - \delta$, pct. per year	7.0	[7, 7]	6.6	7.0	8.3
Capital stock, pct. above the no-AI path	13.5	[6, 27]	2.3	13.8	56.3
Labor share, pct. of income	57.2	[55, 59]	59.4	56.1	45.2
Capital share, pct. of income	42.8	[41, 45]	40.6	43.9	54.8
C. The labor market					
Cognitive employment, pct. change since mid-2026	−4.2	[−9, −1]	−0.5	−3.9	−21.5
Unemployment rate, cognitive workers, pct.	4.6	[3, 8]	2.9	4.5	17.9
Unemployment rate, all workers, pct.	4.6	[4, 6]	3.9	4.6	11.9

Note: We convert each respondent's five answers to the model's parameters as in Section 3.5 and run them through the model. The columns report the median and the 25th and 75th percentiles of each outcome across the 3259 respondents who answered all five items (survey weights). All of the simulations hold the reinstatement ratio, the posting speed, the wage rigidity, and the elasticity of capital supply at the substantial change scenario's values ($\rho = 0.25$, $\theta^H = 0.25$, $\xi = 0.5$, $\varepsilon = 3$).

The public's answers deliver outcomes close to the substantial change scenario. The median American expects a capable AI, one that handles six of the eight tasks by 2030, but expects diffusion to be limited. AI is deployed on two-fifths of the tasks it could perform, it augments rather than automates half the time, and it raises the productivity of tasks impacted by AI by around 55 percent. In the model, the medians of the outcomes across respondents are close to the substantial scenario: GDP is 8.6 percent above its no-AI path and the unemployment rate in both cognitive occupations and in the broader economy is around 4.6 percent.

4.5 Robustness: supply of capital

The elasticity of capital supply is one number behind Table 3 that is neither measured nor a dimension of the scenarios, so as a robustness exercise Table 5 reruns the substantial and extreme change scenarios at $\varepsilon = 1$, 6, and ∞ in place of our $\varepsilon = 3$, holding everything else at its values in the table. Two results stand out.

First, the elasticity moves prices and quantities of capital substantially. When capital has a very low supply elasticity of $\varepsilon = 1$ in the extreme scenario, for example, the net return to capital rises to 10.3 percent versus 6.5 percent when capital is perfectly elastic.

Table 5: Scenarios in 2030 at four elasticities of capital supply.

	Substantial				Extreme			
ε:	1	3	6	∞	1	3	6	∞
GDP, pct. above the no-AI path	6.4	8.3	9.1	10.0	21.3	32.4	37.2	43.3
Average wage, pct. above the no-AI path	−1.6	2.1	3.7	5.6	−9.2	9.7	18.3	30.1
Net return to capital $r_t - \delta$, pct.	7.6	7.0	6.8	6.5	10.3	8.3	7.5	6.5
Capital stock, pct. above the no-AI path	9.3	13.8	15.7	18.2	33.5	56.3	67.1	82.2
Labor share, pct. of income	55.1	56.1	56.5	57.0	41.2	45.2	46.9	49.1

Note: The substantial and extreme change scenarios of Table 3 (the $\varepsilon = 3$ columns) rerun with the capital supply schedule (6) at $\varepsilon = 1$, 6, and ∞, the last being a rental rate pegged at $\bar{r}$; all other parameters as in Table 1.

Second, the wage and the return to capital move in opposite directions, and at $\varepsilon = 1$ the wage changes sign. When capital is inelastic, wages can decline (relative to the no-AI path, and in the extreme scenario even in absolute terms).

For example, if AI is used widely but makes the work it touches only marginally more productive (what Acemoglu and Restrepo (2019) call "so-so technologies") the fall in labor income can be large relative to the increase in GDP. In this case compensating workers would consume most of the gains. The model can produce such an outcome with parameter values inside the ranges we consider in this paper.[14]

4.6 Robustness: wage rigidity

The rigidity parameter ξ governs how fast the cognitive wage moves toward its market-clearing level: the gap between the two shrinks by the factor ξ a year. When $\xi = 0$, wages are fully flexible and adjust to clear the labor market while when $\xi = 0.9$, real wages adjust slowly with a half-life of seven years. Table 6 reruns the substantial and extreme change scenarios at $\xi = 0.75$ and 0.9 in place of 0.5, and the extreme scenario at $\xi = 0$, the flexible wage. The table shows the trade-off the parameter governs: a stickier wage means more layoffs and a smaller wage decline.

In the extreme scenario, the cognitive wage in 2030 is 11.5 percent below its no-AI path at our baseline $\xi = 0.5$ and the cognitive unemployment rate is 17.9 percent. When wages are very rigid ($\xi = 0.9$, half-life of 7 years), the cognitive wage is 2.8 percent above the no-AI path but cognitive unemployment rises to 24.0 percent. On the other side of the spectrum, when wages are fully flexible ($\xi = 0$), the cognitive wage falls by more than 42 percent while cognitive unemployment is just 2.6 percent. The effects of AI are felt in the cognitive labor market in either wages or unemployment; our baseline scenario shares the cost across both margins.

[14] Suppose AI capability and adoption and productivity gain follow the substantial scenario ($m_{2030} = 0.3$, $d_{2030} = 0.4$, $a_{2030} = 0.45$), while the automation share, re-employment time and task reinstatement take their extreme-scenario values ($\psi = 0.9$, $\mu = 0.04$, $\rho = 0$), while the supply of capital is inelastic, $\varepsilon = 1$. In this case, the 2030 GDP is 7.2 percent above its no-AI path, but total labor income is lower by 4.3 percent of no-AI GDP. Holding cognitive occupations' income at its no-AI level would take a transfer equal to 84 percent of GDP gains.

Table 6: Scenarios in 2030 at four rigidities of the cognitive wage.

	Substantial			Extreme			
ξ:	0.5	0.75	0.9	0	0.5	0.75	0.9
GDP, pct. above the no-AI path	8.3	7.9	7.7	36.6	32.4	30.5	29.2
Average wage, pct. above the no-AI path	2.1	2.2	2.3	1.6	9.7	11.1	11.9
Cognitive wage $w_{C,t}$, pct.	−0.3	0.7	1.4	−42.2	−11.5	−2.9	2.8
All-other wage $w_{N,t}$, pct.	5.9	4.5	3.7	70.1	33.6	25.8	21.1
Cognitive employment, pct. since mid-2026	−3.9	−4.6	−5.0	−1.3	−21.5	−25.9	−28.5
Unemployment rate, cognitive workers, pct.	4.5	5.1	5.4	2.6	17.9	21.7	24.0
Unemployment rate, all workers, pct.	4.6	4.9	5.2	3.1	11.9	13.9	15.2

Note: The scenarios of Table 3 rerun with the rigidity of the cognitive wage in Equation (30) at $\xi = 0$, 0.75, and 0.9 in place of 0.5; all other parameters as in Table 1.

5. Conclusions

This paper seeks to place the wide range of views on AI's economic effects inside a common framework, as alternative settings of a small number of parameters that can be reasoned about and, increasingly, measured. The framework is a standard task-based model with two groups of occupations, an ideas channel, and a frictional labor market. The scenario explorer that accompanies the paper lets anyone set the parameters and see the paths they imply for GDP, growth, wages, the return to capital, the labor share, and unemployment through 2030.

The scenarios yield four main results. First, the range of outcomes is wide. In the modest change scenario, AI acts like a "normal technology" over the next five years and has small macroeconomic effects. In the extreme scenario, AI is transformative: by 2030 GDP is 32 percent above its no-AI path and nearly one in five cognitive workers is unemployed. Second, almost all of the divergence comes after 2027, a year and a half from the mid-2026 anchor, because the scenarios share today's readings and separate only as the reach and use of AI diverge. Third, the mechanics are the same in every scenario: output rises, cognitive employment shrinks, and employment in all other occupations grows, and how much unemployment the reallocation creates depends both on its size and on how easily displaced workers find jobs outside their old occupations. What the scenarios disagree about is size and speed.

Finally, the labor market consequences depend on how easily the economy adds capital and on how quickly relative wages adjust. A more elastic supply of capital lets labor keep more of the gain; a less elastic one can push the average wage below its no-AI path. The labor share falls from 60 percent of income to 56 and 45 percent in the substantial and extreme scenarios. Within labor, the cognitive wage falls below its no-AI path in the extreme scenario, while wages in all other occupations, where labor has become relatively scarce, rise sharply. The cost to cognitive workers is a combination of a lower relative wage and higher unemployment. When wages adjust quickly, more of the cost is in wages; when they are rigid, more is in unemployment (Table 6). Both are clearly costly to the workers affected.

Each parameter underpinning our scenario explorer is something we can potentially measure: the share of tasks AI can do, the share of firms using it, the productivity gain per task, how long a displaced worker takes to find a new job. The scenarios therefore make statements about data we will observe over the coming years, and those data will tell us which scenario we are in. Today, none of the three can be ruled out. The lasting contribution, we hope, is the framework itself: a common way to state disagreements about the economic effects of AI as disagreements about a few measurable parameters.

The questions the framework sharpens are questions about transitions. If AI has effects like the modest or substantial scenarios, the reallocation, while costly, is a size the U.S. labor market has absorbed historically. If something closer to the extreme scenario occurs, the disruption is much larger: 18 percent of the cognitive labor force is unemployed, the relative wage of cognitive workers falls immensely, and 15 percent of GDP shifts from payrolls to capital income. The aggregate gain is nearly three times what cognitive workers lose in wages and employment, so the resources to compensate them exist. But whether and how those resources reach the people who bear the cost is not something growth delivers by itself. The mechanisms through which people could share in the gains of a much richer economy — retraining, income support, universal basic capital or universal basic income, and others not yet designed — then become a central question of economic policy.

Which of these worlds we are heading toward may become clearer within a year or two, and preparing for potential disruption seems to us the prudent course. We are developing this framework and the accompanying scenario explorer to inform the debate over how society can best manage the transition to powerful AI.

A. Simulation procedure

The model is simulated on a monthly grid. Rates quoted as fractions per period are converted to continuously compounded rates. For example, a quit fraction $\hat{q}$ enters as $q = -\ln(1-\hat{q})$. At t_0 the AI objects are near the start of their paths, so that the GDP gap they imply is at most a quarter of a percent, and the flow block starts from the steady state (38), evaluated at the normal-times search discount $\bar{\mu}$ in every scenario. The scenario's own μ enters Equations (33) and (35) from t_0 on.

The simulation involves three groups of objects:

- *Fixed parameters*: the elasticity of substitution σ, the base-period labor share s_{L,t_0}, and the elasticity of capital supply ε, with $\bar{r}$ and δ for reporting the net return (Section 2.1.1); the two groups' shares s_{C,t_0} and s_{N,t_0} and base employment ℓ_{C,t_0} and ℓ_{N,t_0} (Sections 2.1.1 and 2.3); the ideas-production parameters λ and ϕ_R, the baseline growth rate g, and the research share of GDP $\iota_{R,t}$ on its trend g_ι (Section 2.2); and the frictional parameters — the quit-rate components q_o^X and q_o^T, the normal pool $\bar{U}$, the matching curvature ι, the posting speed θ^H, and the rigidity of the cognitive wage ξ and search discount μ, which the scenarios set, with its normal-times value $\bar{\mu}$ — together with the steady-state objects they imply through Equation (38) (Section 2.3).
- *Exogenous scenario paths*: m_t, d_t, a_t, ψ_t, and ρ (Section 2.1.2).
- *Endogenous variables*: the ideas stock; the rental-rate gap and the capital stock; the shift $\tilde{\ell}_{N,t}$, the labor share, the wage, output per worker, measured TFP, and the two employment targets; and, for each group, the gaps, quits, separations, openings, search effort, hires, finding rate, employment, and pool, together with the reported aggregates.

Each month, the equations are evaluated in the following order:

1. **Paths.** Evaluate the scenario paths by Equation (8) and its analogue for ψ_t, at t and at $t+1$. The ideas stock $\Delta \ln A_t$ is predetermined, from step 9 of the previous month.
2. **Capital market.** Solve for the rental-rate gap $\Delta \ln r_t$ by Proposition 1: find the root of Equation (18), at which the capital demanded, Equation (17), equals the capital supplied, Equation (6); one equation in one unknown, monotone, solved by bisection (at $\varepsilon = \infty$ the root is zero).
3. **Wage, shares, and targets.** Form $\tilde{\ell}_{N,t}$, the labor share, the wage, and the capital stock by Proposition 1 at $\Delta \ln r_t$; output per worker by the identity (5) and measured TFP by Equation (45); and the two targets by Equation (13), at t and, from the date-$(t+1)$ paths, at $t+1$.
4. **Gaps and the cognitive wage.** Form the all-other group's shortfall against the date-$(t+1)$ target, Equation (28); the attached cognitive force, Equation (29); the wage that would clear it and the prices at current employment, both by the system (39); the sticky wage by Equation (30); and cognitive labor demand at that wage, again by the system (39) solved for $\ell_{C,t}$ given $w_{C,t}$.

5. **Separations and openings.** Set the quit rates by Equation (27) from last month's finding rates, layoffs by Equation (31), and openings by Equation (32).

6. **Matching.** Form effective search by Equation (33), hires by Equation (34), and the finding rates by origin by Equation (35).

7. **Stocks.** Update employment and the pools by Equations (36) and (37).

8. **Reporting.** Record the actual economy's GDP, all-other wage, return, capital, and labor share from the system (39) at realized employment, with the cognitive wage paid and the average wage of the employed; the pool in excess of its normal level, the reallocation flow, and the aggregate overhang (Table A.1, panel D); and the cognitive shift $\tilde{\ell}_{C,t} = \ln(\ell^*_{C,t}/\ell_{C,t_0})$.

9. **Ideas.** Set the uplift $\Delta \ln R_t$ to the GDP gap in Equation (22), and update $\Delta \ln A_{t+1}$ by a step of length $h = 1/12$ on Equation (42).

Table A.1 displays the full model in the form in which it is evaluated: 44 equations in 44 unknowns per month, plus the steady state solved once at the start. Where a variable has both an exact and a first-order form the table gives both, the exact one on the variable's own row and the first-order one, marked ≈, on the row beneath it; the simulation uses the exact equations; the approximations are provided only for intuition.

With the exact set, two rows of the table are approximate: measured TFP, which is the base-price index, and the ideas update, which is a monthly step on an exact recursion (on the extreme path it is within 0.02 percentage points of the closed form in 2030). Every other row holds exactly given its right-hand side. Apart from the rental-rate row and the actual economy, every right-hand side contains only parameters, exogenous paths, variables predetermined at t ($\Delta \ln A_t$, $\ell_{o,t}$, $U_{o,t}$, $f_{o,t-1}$, and $w_{C,t-1}/w_{t-1}$), or variables determined in an earlier row, so the model is solved by recursion; the ideas rows of panel A use the GDP of panel D and deliver next month's stock. The rental-rate row is one monotone equation in one unknown each month, solved by bisection, and the actual economy is two nested ones. Table A.2 lists the parameters and exogenous objects.

Table A.1: The two-group system written out: the equation determining each endogenous variable, by block.

Endogenous variable	Eq.	Equation
A. Ideas (3 equations)		
Research uplift, $\Delta \ln R_t$	(22)	$\Delta \ln R_t = \Delta \ln Y_t$, with Y_t the actual economy's GDP (panel D)
Growth gap, Δg_t	(42)	$\Delta g_t = g\left[e^{\lambda \Delta \ln R_t - (1-\phi_R)\Delta \ln A_t} - 1\right]$
Ideas stock, $\Delta \ln A_{t+1}$	(42), (43)	$\Delta \ln A_{t+1} \approx \Delta \ln A_t + h\,\Delta g_t$
B. Capital market, wage, shares, and targets (9 equations)		
Rental gap, $\Delta \ln r_t$	(18), (6)	$\varepsilon \Delta \ln r_t = \Delta \ln K_t$, the capital row evaluated at $\Delta \ln r_t$; one root, zero at $\varepsilon = \infty$
first order	(11), (44)	$\Delta \ln r_t \approx \left[m_t d_t a_t + ((1-\rho) - (1-\sigma)a_t)\psi_t m_t d_t / s_{K,t_0} + \Delta \ln A_t\right] / (\varepsilon + \sigma / s_{L,t_0})$
Shift of N's demand, $\tilde{\ell}_{N,t}$	(15)	$\tilde{\ell}_{N,t} = -\ln\left(1 - m_t d_t \left[1 - \rho\psi_t - (1-\psi_t)\, e^{-(1-\sigma)a_t}\right]\right)$
first order	(19)	$\tilde{\ell}_{N,t} \approx (1-\rho)\,\psi_t m_t d_t + (1-\sigma)(1-\psi_t)\, m_t d_t a_t$
Labor share, $s_{L,t}$	(14), (44)	$s_{L,t} = 1 - \left[s_{K,t_0} + s_{L,t_0}\psi_t m_t d_t \left(e^{-(1-\sigma)a_t} - \rho\right)\right] e^{(1-\sigma)\Delta \ln r_t}$; $\Delta \ln s_{L,t} = \ln(s_{L,t}/s_{L,t_0})$
first order	(11), (44)	$\Delta \ln s_{L,t} \approx -(1-\rho)\,\psi_t m_t d_t + (1-\sigma)\left[\psi_t m_t d_t a_t - \frac{s_{K,t_0}}{s_{L,t_0}} \Delta \ln r_t\right]$
Wage, $\Delta \ln w_t$	(16), (44)	$\Delta \ln w_t = \left(\Delta \ln s_{L,t} + \tilde{\ell}_{N,t}\right)/(1-\sigma) + \Delta \ln A_t$
first order	(11), (44)	$\Delta \ln w_t \approx m_t\, d_t\, a_t + \Delta \ln A_t - (s_{K,t_0}/s_{L,t_0})\,\Delta \ln r_t$
Output per worker, $\Delta \ln(Y_t/L_t)$	(5)	$\Delta \ln(Y_t/L_t) = \Delta \ln w_t - \Delta \ln s_{L,t}$
Capital, $\Delta \ln K_t$	(17)	$\Delta \ln K_t = \ln\left((1-s_{L,t})/s_{K,t_0}\right) + \Delta \ln(Y_t/L_t) - \Delta \ln r_t$
Measured TFP, $\Delta \ln \mathrm{TFP}_t$	(45)	$\Delta \ln \mathrm{TFP}_t \approx -\frac{1}{1-\sigma} \ln\left[s_{K,t_0} + s_{L,t_0}\left(1 - m_t d_t \left(1 - e^{-(1-\sigma)a_t}\right)\right) e^{-(1-\sigma)\Delta \ln A_t}\right]$
first order	(26)	$\Delta \ln \mathrm{TFP}_t \approx s_{L,t_0}\left(\Delta \ln A_t + m_t d_t a_t\right)$
Target of N, $\ell^*_{N,t}$	(13)	$\ell^*_{N,t} = \ell_{N,t_0}\, e^{\tilde{\ell}_{N,t}}$
Target of C, $\ell^*_{C,t}$	(13)	$\ell^*_{C,t} = \ell_{C,t_0} + \ell_{N,t_0} - \ell^*_{N,t}$
C. Flows, $o, j \in \{C, N\}$ (23 equations: ten pairs and three for the cognitive group)		
Overhang, $G_{o,t}$	(28)	$G_{o,t} = \max\{0,\ \ln \ell_{o,t} - \ln \ell^*_{o,t+1}\}$
Shortfall, $B_{o,t}$	(28)	$B_{o,t} = \max\{0,\ \ln \ell^*_{o,t+1} - \ln \ell_{o,t}\}$
Quit rate, $q_{o,t}$	(27)	$q_{o,t} = q_o^X + q_o^T\, f_{o,t-1}/\bar{f}_o$
Attached cognitive force, $N_{C,t}$	(29)	$N_{C,t} = \ell_{C,t} + \max\{0,\ U_{C,t} - \bar{U}_C\}$
Sticky cognitive wage, $w_{C,t}$	(30)	$\ln(w_{C,t}/w_t) = \xi^{1/12} \ln(w_{C,t-1}/w_{t-1}) + (1-\xi^{1/12}) \ln(w^c_{C,t}/w_t)$; $w^c_{C,t}$ clears $N_{C,t}$ by (39)
Cognitive demand, $\ell^d_{C,t}$	(39)	the system (39) solved for $\ell_{C,t}$ at $w_{C,t}$ and $\ell_{N,t}$; $E_t = \max\{0, \ell_{C,t} - \ell^d_{C,t}\}$, $Z_t = \max\{0, \ell^d_{C,t} - \ell_{C,t}\}$
Layoffs, $D_{C,t}$	(31)	$D_{C,t} = \max\{0,\ E_t - q_{C,t}\,\ell_{C,t}\}$; $D_{N,t} \equiv 0$
Openings, $v_{o,t}$	(32)	$v_{C,t} = \bar{\pi}_C^{-1}\left[\max\{0, q_{C,t}\ell_{C,t} - E_t\} + \theta^H Z_t\right]$; $v_{N,t} = \bar{\pi}_N^{-1}\left[(q_{N,t} + \theta^H B_{N,t})\,\ell_{N,t}\right]$
Effective search, $S_{j,t}$	(33)	$S_{C,t} = U_{C,t} + \mu\, U_{N,t}$, $S_{N,t} = \mu\, U_{C,t} + U_{N,t}$
Hires, $H_{j,t}$	(34)	$H_{j,t} = \chi\, S_{j,t}\, v_{j,t} / \left(S^\iota_{j,t} + v^\iota_{j,t}\right)^{1/\iota}$
Finding rate by origin, $f_{o,t}$	(35)	$f_{C,t} = H_{C,t}/S_{C,t} + \mu\, H_{N,t}/S_{N,t}$, $f_{N,t} = \mu\, H_{C,t}/S_{C,t} + H_{N,t}/S_{N,t}$
Employment, $\ell_{o,t+1}$	(36)	$\ell_{o,t+1} = (1-q_{o,t})\,\ell_{o,t} - D_{o,t} + H_{o,t}$
Pool by origin, $U_{o,t+1}$	(37)	$U_{o,t+1} = U_{o,t} + q_{o,t}\,\ell_{o,t} + D_{o,t} - f_{o,t}\,U_{o,t}$
D. Reporting (9 equations, the actual economy counting five)		
Cognitive shift, $\tilde{\ell}_{C,t}$	(19)	$\tilde{\ell}_{C,t} = \ln\left(\ell^*_{C,t}/\ell_{C,t_0}\right)$ $\left(\approx -(s_{N,t_0}/s_{C,t_0})\,\tilde{\ell}_{N,t}\right)$

continued on the next page

Table A.1, continued

Endogenous variable	Eq.	Equation
Excess unemployment, u_t^x	—	$u_t^x = (U_{C,t} + U_{N,t} - \bar{U})/L$, the pool above its normal level
Actual economy, Y_t, $w_{N,t}$, r_t, K_t, $s_{L,t}$	(39)	the price index, the two groups' labor demands, and the supply of capital at realized employment $(\ell_{C,t}, \ell_{N,t})$, $$\begin{aligned} & s_{L,t_0}\Lambda_{C,t}\, e^{(1-\sigma)\Delta \ln w_{C,t}} + s_{N,t_0}\, e^{(1-\sigma)\Delta \ln w_{N,t}} + B_t\, e^{(1-\sigma)\Delta \ln r_t} = 1, \\ & \frac{\ell_{C,t}}{\ell_{C,t_0}} = \frac{\Lambda_{C,t}}{s_{C,t_0}/s_{L,t_0}}\, e^{\Delta \ln(Y_t/\bar{L}) - \sigma \Delta \ln w_{C,t}}, \qquad \frac{\ell_{N,t}}{\ell_{N,t_0}} = e^{\Delta \ln(Y_t/\bar{L}) - \sigma \Delta \ln w_{N,t}}, \\ & \Delta \ln K_t = \varepsilon\, \Delta \ln r_t, \end{aligned} \tag{39}$$ with $\Lambda_{C,t} \equiv s_{C,t_0}/s_{L,t_0} - m_t d_t\,[1 - \rho\psi_t - (1-\psi_t)e^{-(1-\sigma)a_t}]$ the surviving mass of cognitive instances, B_t the bracket of the labor-share row, $\bar{L} = L - \bar{U}$, and the wages deflated by A_t; the cognitive wage is the sticky wage $w_{C,t}$, or $\mathrm{MPL}_{C,t}$ where employment trails demand. Evaluated at $(N_{C,t}, \ell_{N,t})$ the system gives $w^c_{C,t}$; solved for $\ell_{C,t}$ at $w_{C,t}$ it gives $\ell^d_{C,t}$; at the targets it is Proposition 1. Average wage $\bar{w}_t = (w_{C,t}\ell_{C,t} + w_{N,t}\ell_{N,t})/(\ell_{C,t} + \ell_{N,t})$; cognitive firms' profit $(\mathrm{MPL}_{C,t} - w_{C,t})\,\ell_{C,t}$, at most half a percent of GDP on our paths
Reallocation, X_t	—	$X_t = \frac{1}{2L}\left(\lvert\ell_{C,t+1} - \ell_{C,t}\rvert + \lvert\ell_{N,t+1} - \ell_{N,t}\rvert\right)$
Aggregate overhang, G_t	—	$G_t = G_{C,t}\,\ell_{C,t}/L$
Per month: 44 equations in the 44 unknowns of panels A–D, of which two are approximate when the exact rows are used and six when the first-order rows are.		
E. Steady state, solved once at t_0 *(9 objects; not part of the monthly count)*		
Hires, pools, $\bar{H}_o$, $\bar{U}_o$	(38)	$\bar{H}_o = \bar{q}_o\,\ell_{o,t_0} = \bar{f}_o\,\bar{U}_o$, $\quad \bar{U}_C + \bar{U}_N = \bar{U}$
Finding rates, $\bar{f}_o$	(38)	$\bar{f}_C = \bar{H}_C/\bar{S}_C + \bar{\mu}\,\bar{H}_N/\bar{S}_N$, $\quad \bar{f}_N = \bar{\mu}\,\bar{H}_C/\bar{S}_C + \bar{H}_N/\bar{S}_N$
Filling rates, $\bar{\pi}_o$, χ	(38)	$\bar{\pi}_o = \chi\left[1 - (\bar{H}_o/\chi\bar{S}_o)^\iota\right]^{1/\iota}$, $\quad (\ell_{C,t_0}\bar{\pi}_C + \ell_{N,t_0}\bar{\pi}_N)/(L - \bar{U}) = 0.65$

Notes. Rows written with = hold exactly given their right-hand sides; rows written with ≈ hold to first order. The exact rows of panel B read the technology instance by instance (Proposition 1); "first order" rows are the alternatives to the row above them, and one set or the other is run. $h = 1/12$; $s_{K,t_0} = 1 - s_{L,t_0}$; $\bar{q}_o = q_o^X + q_o^T$; $\bar{S}_j$ is the effective-search row at the steady-state pools and the normal-times discount $\bar{\mu}$. Rows in o or j stand for one equation per group; on the scenario paths $G_{N,t}$ is zero and $B_{C,t}$ is negligible, so that layoffs occur only in C and shortfall postings only in N, as in the text. The targets at $t+1$ in panel C are the $\tilde{\ell}_{N,t}$ and target rows of panel B evaluated at the date-$(t+1)$ paths, which are exogenous; they do not involve the rental rate. Predetermined at t are $\Delta \ln A_t$, $\ell_{o,t}$, $U_{o,t}$, $f_{o,t-1}$, and $w_{C,t-1}/w_{t-1}$, with initial conditions $\Delta \ln A_{t_0} = 0$, $\ell_{o,t_0} = (s_{o,t_0}/s_{L,t_0})(L - \bar{U})$, $U_{o,t_0} = \bar{U}_o$, $f_{o,t_0-1} = \bar{f}_o$, and $w_{C,t_0-1} = w_{t_0-1}$. Panel E states ten conditions for nine objects; one is redundant, because the finding-rate conditions determine the pools only up to scale, which the adding-up condition fixes.

Table A.2: Parameters and exogenous objects taken as given by the system in Table A.1.

Symbol	Meaning (substantial scenario's value where the text gives one)
A. Technology and factor markets (Section 2.1.1)	
σ	elasticity of substitution across tasks (0.5)
ω_i, $m_i = \omega_i/s_{L,t_0}$	base-period expenditure share (Domar weight) of task i, and its mass, the same weight relative to the pre-AI wage bill; each sums to one
s_{L,t_0}, $s_{K,t_0} = 1 - s_{L,t_0}$	base-period labor and capital shares (0.6, 0.4)
L	labor force; head counts are shares of L_t, which grows at n, so L is constant in these units

continued on the next page

Table A.2, continued

Symbol	Meaning (substantial scenario's value where the text gives one)
ε	elasticity of the capital stock with respect to the rental rate, $\Delta \ln K_t = \varepsilon \Delta \ln r_t$ (3; ∞ is a rental rate pegged by world capital markets)
$\bar{r}$, δ	no-AI gross rental rate and depreciation, used to report the net return $r_t - \delta$ (0.115, 0.05: 6.5 percent without AI)
t_0, h	start of the AI gaps (2024) and period length in years (1/12)
B. The two groups of occupations (Sections 2.1.1 and 2.3)	
s_{C,t_0}, s_{N,t_0}	base-period wage-bill shares of the cognitive and of all other occupations, $s_{C,t_0} + s_{N,t_0} = s_{L,t_0}$; the affected tasks must fit inside C, $\bar{m} \le s_{C,t_0}/s_{L,t_0}$
$s_{C,t_0}/s_{L,t_0}$, $s_{N,t_0}/s_{L,t_0}$	the groups' shares of employment (0.624 and 0.376, CPS 2025, so $s_{N,t_0}/s_{C,t_0} \approx 0.6$)
ℓ_{C,t_0}, ℓ_{N,t_0}	base employment, $(s_{o,t_0}/s_{L,t_0})(L - \bar{U})$ for $o = C, N$; the two add up to $L - \bar{U}$
C. Ideas production (Section 2.2)	
λ, ϕ_R	returns to research input and to the ideas stock in ideas production ($\lambda = 1$; $1 - \phi_R = s_{L,t_0}(1 - \phi) + \lambda = 2.86$ in labor-augmenting units with research input in goods, from $1 - \phi = 3.1$ in TFP units and researcher-equivalents)
g; $g_A = s_{L,t_0}\, g$	no-AI growth of the labor-augmenting ideas stock (0.0167 per year) and of measured TFP (0.01)
n	growth of the labor force (0.0033 per year); it sets GDP growth without AI at $g + n = 0.02$ and enters only the levels
$\iota_{R,t}$; g_ι	research share of GDP, funded by lump-sum taxes (0.035 in 2024), and its trend ($g_\iota = g_R - g - n \approx 0.028$ per year, with or without AI); only the gap $\Delta \ln \iota_{R,t} = 0$ enters
$g_R = (1 - \phi_R)\, g/\lambda$	implied baseline growth of research input in goods (≈ 0.048 per year; 0.031 in researcher-equivalents); derived
$\gamma = \lambda/(1 - \phi_R)$	long-run elasticity of the ideas stock with respect to research input (≈ 0.35; 0.21 in measured TFP); derived
ν	scale of ideas production; drops out of the gaps
D. Job-flow parameters (Section 2.3)	
q_o^X, q_o^T	exogenous and finding-rate-sensitive components of the quit rate, $o \in \{C, N\}$
$\bar{q}_o = q_o^X + q_o^T$	normal quit rate (economy-wide $\bar{q} = 0.11$ per year, of which 0.06 stands for occupational moves and 0.05 for exits from the labor force)
θ^H	fraction of the shortfall posted per month (set by the scenario: 0.25; modest 0.10, extreme 0.50)
ξ	rigidity of the cognitive wage per year, the weight on last month's gap to the common wage at the monthly step $\xi^{1/12}$ (0.5 in every scenario)
μ; $\bar{\mu}$	search discount of a worker searching outside the group of origin: on the path, set by the scenario (0.08; modest 0.17, extreme 0.04); in the steady state, $\bar{\mu} = 0.17$ in every scenario
ι, χ	matching curvature (1.27) and efficiency, set so that the mean filling rate is 0.65 per month
$\bar{U}$; $\bar{f}$	normal pool, taken from the data (3.8 percent of L, CPS 2025); implied aggregate finding rate $\bar{q}\,(L - \bar{U})/\bar{U}$
$\bar{U}_o$, $\bar{f}_o$, $\bar{\pi}_o$, $\bar{H}_o$	steady-state pools, finding and filling rates, and hires by group (Table A.1, panel E)
E. Scenario paths for goods and services (Section 2.1.2); all of them fall on the cognitive occupations	
m_t	affected mass: fraction of the economy's tasks that AI affects, at pre-AI wage-bill weights (0.14 in mid-2026); the affected fraction of C's own tasks is $(s_{L,t_0}/s_{C,t_0})\, m_t$
d_t	diffusion: share of affected instances performed with AI (0.10)
a_t	log productivity gain, or unit-cost decline, per AI-performed instance (0.35)

continued on the next page

Table A.2, continued

Symbol	Meaning (substantial scenario's value where the text gives one)
ψ_t	automation share: fraction of AI-performed instances that are automated (0.75; modest 0.50, extreme 0.90; constant in the scenarios)
ρ	reinstatement ratio: new labor-task mass per unit of automated mass, the new tasks belonging to C (0.25; modest 0.50, extreme 0; constant)
$\bar{m}, \kappa_m, t_m; \bar{d}, \kappa_d, t_d$	ceilings, slopes, and midpoints of the logistic paths of m_t and d_t; the ceilings are the natural maxima, $\bar{m} = s_{C,t_0}/s_{L,t_0}$ and $\bar{d} = 1$; a scenario sets the values at the start of 2030 ($m_{2030} = 0.3$, $d_{2030} = 0.4$; modest 0.2, 0.2; extreme 0.5, 0.6), and slope and midpoint follow from them and the mid-2026 anchors by Equation (8′) ($\kappa_m = 0.33$, $\kappa_d = 0.51$; modest 0.14, 0.23; extreme 0.75, 0.74)
a_0, g_a	level of the gain at t_0 and its linear slope per year ($g_a = 0.028$; modest 0, extreme 0.10; the gain in mid-2026 is 0.35, 0.30, and 0.45)
$\underline{\psi}, \bar{\psi}, \kappa_\psi, t_\psi$	floor, ceiling, slope, and midpoint of the logistic path of ψ_t

B. The survey of US adults

This appendix documents the survey reported in Section 3.5. We converted every respondent's answers to implied values of m_{2030}, d_{2030}, ψ, a_{2030}, and μ. Table 2 reports the medians and quartiles of those values, interpolated within the answer bins as described at the end of this appendix, and Table 4 reports the quantiles, across respondents, of the outcomes the model implies at each respondent's values.

B.1 Samples

US adults. Morning Consult fielded the questionnaire to its online panel of US adults in eight daily samples of about 2,000, August 11–23, 2026. Some respondents answered more than once and we kept each respondent's first response. The sample contains 10,980 unique people whose responses are weighted to the adult population on age, gender, education, race, region, and income. Morning Consult screens responses for speed and attention (a minimum timer, an unlikely-event screener, a semantic-association check). "Not sure" answers are excluded item by item, so the number of respondents on each question differs. The questionnaire was developed in partnership with Morning Consult over three pilots in July and August 2026, each with about 2,000 respondents. The respondents were asked about their expectations for either 2027 or 2030. The year was chosen at random and used throughout a respondent's questionnaire. Figure 1 depicts results for both years, but when estimating the implied values of the parameters, only the 2030 data is used, since the scenarios are specified by their 2030 values.

B.2 Questionnaire

In the following, we describe the questions posed in the survey:

1. **Capability.** *When do you think AI software will be able to do each of the following as well as*

a skilled professional? Focus on capability — regardless of whether people or companies will actually use AI to do it.

Eight tasks are listed in a grid. Rows are randomly shuffled. For each, the respondent picks one of: already can, by 2027, by 2030, later than 2030, never, not sure. The eight tasks, in the order respondents rank their difficulty, are:

1. Write routine business emails and documents
2. Do a small business's monthly bookkeeping
3. Prepare an accurate personal tax return from someone's documents
4. Draft and negotiate a routine commercial contract
5. Build, ship and maintain a working software product
6. Design, run and write up a scientific study that gets published
7. Run a small online business end to end profitably, with no human employees
8. Independently make a scientific discovery important enough to win a Nobel Prize

2. **Adoption.** *Even when AI is technically capable of performing a task, businesses and people don't always choose to use it that way. Of the tasks AI is capable of performing, what share do you think people will regularly use it for at work in 2030?*

   We offered five bins: less than 10 percent, 10–24, 25–49, 50–74, more than 75 percent.

3. **Automation.** *Now imagine it is 2030 and AI is useful for each of the following tasks. For each task below, do you think AI will have taken over the task from workers, or will workers be doing it with AI's help? Workplaces will differ — answer for the arrangement you expect to be most common in 2030.*

   Five of the above tasks (emails, tax return, bookkeeping, contract, software product) are listed in a grid with rows randomly shuffled. For each task the respondent chooses one of: (a) AI does it alone, (b) AI does most of it, with workers checking, (c) workers do most of it, with AI helping, (d) workers do it alone.

4. **Productivity gain.** *From your own experience: think about a task at your work that today's AI is well suited for, like summarizing a document or drafting an email. How long do you think it takes on average to finish the task using AI — counting the time spent checking and fixing the AI's work — compared with not using AI?*

   Six options: (a) it takes longer, (b) about the same time, (c) a little less time, (d) about half the time, (e) about a quarter of the time, (f) about a tenth of the time or less.

   Respondents who have not used AI at work are asked the same question about "a worker" they have seen or heard of.

5. **Re-employment.** *In normal times, someone who loses a job typically takes about 3 months to find a new one. Now think about a person who loses their job in 2030 due to AI automation and needs to switch into a different occupation. How long do you think it would take them to find a new job?*

Eight bins: (a) less than 1 month, (b) 1–2 months, (c) 3–4, (d) 5–6, (e) 7–12, (f) 13–24, (g) 2–3 years, (h) more than 3 years, or would never find a new job.

B.3 Mapping answers to parameters

Table B.1 describes how we map each survey respondent's answers to parameter values in the scenario explorer.

Capability. The model's m_{2030} is a share of the wage bill while in the survey we offer eight specific tasks. The tasks were chosen to systematically represent increasing levels of difficulty from routine office work to Nobel-level frontier science and to be recognizable to a general audience. To map a respondent's survey response to the parameter, we count the number of the eight tasks they expect AI to do by 2030, as a share of the rows they answered, multiplied by 95 percent to capture that a fraction of knowledge work will continue to be done in person. Finally, we multiply by the knowledge-work share of the wage bill (0.62). For example, seven tasks out of eight corresponds to $m_{2030} = 7/8 \times .95 \times .62 = .52$. The three scenarios' values of m_{2030} correspond to about three, four, and seven tasks performed by AI.

The adult survey also asks employed respondents what share of their own working hours goes to tasks AI will be able to do about as well as they can by 2030. Among workers in cognitive occupations the median answer is 27 percent, against 71 percent from the task grid: respondents expect more of the listed tasks to be within AI's reach than of their own work, a pattern common in surveys about new technology.

Adoption. Adults answered in five bins and we code a respondent's answer as the bin's midpoint, so "25–49 percent" is $d_{2030} = 0.37$. When we need aggregate statistics we use interpolated values, as described below.

Automation. We set ψ so it corresponds to the share of the five listed tasks that AI can perform on its own. We code the answers that respondents could select from as follows: "AI does it alone" counts as 1, "AI does most of it, with workers checking" counts as $\frac{1}{2}$, and "workers do most of it, with AI helping" counts as 0 automation. A row where the respondent selects "workers do it alone" is a task AI does not perform and is dropped from both numerator and denominator. As an example, a respondent who marks two rows as "AI does it alone," two "AI does most of it," and one "workers do most of it" has an automation score of $\psi = (2 + 2 \times \frac{1}{2})/5 = 0.6$. Respondents who answer "not sure" on more than two rows are left out.

Productivity gain. The survey asks how long AI takes for a task compared to a human without AI, on a task AI is well suited for, counting the time spent checking and fixing the AI's work. Each option is read as a time ratio — "about a quarter of the time" is 0.25 — and the gain is the log of the speed-up, $a_{2030} = \ln(1/0.25) = 1.39$. "It takes longer" and "about the same time" are both coded as a gain of zero, as productivity losses could be avoided by not using AI.

Table B.1: Every survey answer and the value it is coded to.

Answer	Coded value	
Capability, m_{2030}: number of the eight tasks marked "by 2030" or earlier		
none	0% of knowledge work	$m_{2030} = 0$
one	12%	0.07
two	24%	0.15
three	36%	0.22
four	48%	0.30
five	59%	0.37
six	71%	0.44
seven	83%	0.52
all eight	95%	0.59
Adoption, d_{2030}: share of capable tasks AI will be used for		
less than 10%		$d_{2030} = 0.05$
10–24%		0.17
25–49%		0.37
50–74%		0.62
more than 75%		0.87
Automation, ψ: each of five tasks		
AI does it alone	counts 1	
AI does most of it, with workers checking	counts $\frac{1}{2}$	
workers do most of it, with AI helping	counts 0	
workers do it alone	task not performed by AI; left out	
		ψ = sum of the counts / tasks AI performs
Productivity gain, a_{2030}: time with AI compared with without		
it takes longer		$a_{2030} = 0$
about the same time		0
a little less time (0.8 of the time)		0.22
about half the time		0.69
about a quarter of the time		1.39
about a tenth of the time or less		2.30
Re-employment, μ: months to a job in a new occupation		
less than 1 month	0.5 months	$\mu = 1$
1–2 months	1.5	0.34
3–4 months	3.5	0.15
5–6 months	5.5	0.09
7–12 months	9.5	0.05
13–24 months	18.5	0.03
2–3 years	30	0.02
more than 3 years, or never	never	0

Note: The middle column is the intermediate reading where there is one (the task's place on the knowledge-work scale; the count a task row contributes; the bin's midpoint in months); the last column is the model parameter. Capability values are the knowledge-work share times $s_{C,t_0}/s_{L,t_0} = 0.62$. Re-employment values are $0.17{\times}3/\text{months}$, capped at 1.

Re-employment. The survey tells respondents that in normal times a job search takes about three months and asks how long a worker displaced by AI in 2030 would take to find a job in a new occupation. The answer, in months, is converted to the model's search discount by $\mu = 0.17 \times 3/\text{months}$. Three months, the normal-times answer, gives $\mu = 0.17$, the value at which the normal-times steady state is solved in every scenario. "More than 3 years, or would never find a new job" is read as never, $\mu = 0$.

Medians of binned answers. Most items offer a handful of bins, so a plain median is whichever bin the fiftieth percentile falls in, and it jumps a whole bin when a few respondents cross an edge. In the adult sample the fiftieth percentile sits almost exactly on the edge between "a little less time" and "about half the time." We therefore report interpolated medians and quartiles, which treat respondents as spread evenly within their bin: the cumulative distribution is linear inside each bin, and the quantile is where it reaches one half (one quarter, three quarters). For adoption and re-employment time the bins are taken directly from the questionnaires. For the time ratio and for the automation parameter (ψ), whose codes are single values, a bin reaches halfway to the neighboring code. The count of tasks is not binned and keeps the plain quantile.

From answers to outcomes. We run the model on each respondent's five values and report the quantiles of each outcome across respondents. This requires all five answers from the same person, so Table 4 uses the 3259 respondents who answered all five items, while the medians of Table 2 use everyone who answered each item. We take those medians item by item, so no single respondent need give all five median answers. Respondents were not asked to forecast the outcomes in Table 4. Those are the model's implications of their answers. For parameters we did not ask about in the surveys — the reinstatement ratio, the posting speed, the wage rigidity, and the elasticity of capital supply — we use the values from the substantial change scenario.

C. The innovation block: derivations

C.1 Fishing out and the research share without AI

Without AI, the ideas stock grows at $g = 1.67$ percent a year and measured TFP at $g_A = s_{L,t_0}\, g = 1.0$ percent, the balanced path of Section 2.1.1. Bloom et al. (2020) measure idea output in units of TFP and research input as R&D spending deflated by the wage of skilled labor, that is, in researcher-equivalents. In those units, they estimate fishing out of $1 - \phi = 3.1$. Two conversions translate their estimate to the units of Equation (20). Measuring ideas by their effect on labor's productivity rather than on TFP multiplies the exponent by s_{L,t_0}, since $\ln A_t$ is $1/s_{L,t_0}$ times log TFP on the balanced path. And a unit of output buys fewer researcher-equivalents as the wage grows, so measuring research input in goods adds λ:

$$1 - \phi_R \;=\; s_{L,t_0}\,(1 - \phi) + \lambda \;=\; 1.86 + 1 \;\approx\; 2.86. \tag{40}$$

By Equation (20), keeping the growth of A_t at g requires research input to grow at $g_R = (1 - \phi_R)\, g/\lambda \approx 4.8$ percent a year, whereas GDP grows at $g + n = 2$ percent. The research share of GDP must therefore rise over time, at

$$g_\iota \;=\; g_R - (g+n) \;=\; \frac{(1-\phi_R)\, g}{\lambda} - (g+n) \;\approx\; 2.8 \text{ percent a year}, \tag{41}$$

which takes $\iota_{R,t}$ from 3.5 percent of GDP in 2024, the share the national accounts record as research and development, to 4.1 percent in 2030. In researcher-equivalents this is growth of research effort at $g_\iota + n = 3.1$ percent a year, in line with its postwar growth. We assume $\iota_{R,t}$ follows this path whether or not AI arrives.

C.2 Cumulation and drag along a scenario path

Equation (24) in the text gives the long-run effect of a single permanent uplift in research input. In the scenarios, a new uplift arrives every year. Each year therefore repeats the impact arithmetic of Equation (23) with that year's $\Delta \ln R_t$, less a fishing-out drag on the level effect $\Delta \ln A_t$ accumulated so far: through the term $A^{\phi_R - 1}$, the ideas already found make further ideas harder to find. Log-differencing Equation (20) against the no-AI path gives $\ln(g_{A,t}/g) = \lambda\, \Delta \ln R_t - (1-\phi_R)\, \Delta \ln A_t$, so the gap in the growth rate is

$$\Delta g_t \;=\; g\left[e^{\lambda\, \Delta \ln R_t - (1-\phi_R)\, \Delta \ln A_t} - 1\right] \;\approx\; g\left[\lambda\, \Delta \ln R_t - (1-\phi_R)\, \Delta \ln A_t\right]. \tag{42}$$

Because Equation (20) makes $A_t^{1-\phi_R}$ linear in cumulated research effort, the level effect also has a closed form, in which the weights embody the baseline research trend $g_R = (1-\phi_R)\, g/\lambda$:

$$\Delta \ln A_t \;=\; \frac{1}{1-\phi_R}\, \ln\!\left(e^{-(1-\phi_R)\, g\, (t-t_0)} \;+\; (1-\phi_R)\, g \int_{t_0}^{t} e^{-(1-\phi_R)\, g\, (t-s)}\, e^{\lambda\, \Delta \ln R_s}\, ds\right). \tag{43}$$

C.3 Pass-through to wages and GDP

The ideas stock enters the factor-price frontier of Section 2.1.3 as a fall in the cost of every labor instance: on top of the level channel's gain, $s_{L,t_0}\, d \ln w_t + s_{K,t_0}\, d \ln r_t = s_{L,t_0}\, d \ln A_t$. Denoting the contributions of the ideas channel by a superscript A,

$$\begin{aligned}
\Delta \ln w_t^A &\approx \Delta \ln A_t - \frac{s_{K,t_0}}{s_{L,t_0}}\, \Delta \ln r_t^A, & \Delta \ln s_{L,t}^A &\approx -(1-\sigma)\, \frac{s_{K,t_0}}{s_{L,t_0}}\, \Delta \ln r_t^A, \\
\Delta \ln (Y_t/L_t)^A &\approx \Delta \ln A_t - \sigma\, \frac{s_{K,t_0}}{s_{L,t_0}}\, \Delta \ln r_t^A, & \Delta \ln K_t^A &\approx \varepsilon\, \Delta \ln r_t^A,
\end{aligned} \tag{44}$$

and these terms add to the level-channel terms of Equation (11).

C.4 Measured TFP beyond first order

Beyond first order, the base-weight counterpart of Equation (26) is the CES price index over all tasks at base factor prices, which under the reading of Proposition 1 is

$$\Delta \ln \mathrm{TFP}_t \;\approx\; -\frac{1}{1-\sigma}\,\ln\Big[\, s_{K,t_0} + s_{L,t_0}\big(1 - m_t\, d_t\big(1 - e^{-(1-\sigma)\,a_t}\big)\big)\, e^{-(1-\sigma)\,\Delta\ln A_t}\Big], \tag{45}$$

with $\sum_i m_i d_{i,t}(1-e^{-(1-\sigma)a_{i,t}})$ in place of the product when tasks differ. Equation (45) is Equation (26) with the curvature of the gains kept, and it agrees with the dual (25) to first order. Equation (45) is the definition of the measured TFP gap used in the simulations.

References

Abraham, Katharine, John Haltiwanger, and Lee Rendell, "How Tight Is the US Labor Market?," *Brookings Papers on Economic Activity*, 2020, pp. 97–163. Spring.

Acemoglu, Daron, "The simple macroeconomics of AI," *Economic Policy*, 2025, *40* (121), 13–58. NBER Working Paper 32487, 2024.

Acemoglu, Daron and David Autor, "Skills, tasks and technologies: Implications for employment and earnings," in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Vol. 4B, Elsevier, 2011, chapter 12, pp. 1043–1171.

Acemoglu, Daron and Pascual Restrepo, "The race between man and machine: Implications of technology for growth, factor shares, and employment," *American Economic Review*, 2018, *108* (6), 1488–1542.

Acemoglu, Daron and Pascual Restrepo, "Automation and new tasks: How technology displaces and reinstates labor," *Journal of Economic Perspectives*, 2019, *33* (2), 3–30.

Acemoglu, Daron and Pascual Restrepo, "Tasks, automation, and the rise in U.S. wage inequality," *Econometrica*, 2022, *90* (5), 1973–2016.

Aghion, Philippe and Simon Bunel, "AI and growth: Where do we stand?," 2024. Working paper.

Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones, "Artificial intelligence and economic growth," in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *The Economics of Artificial Intelligence: An Agenda*, University of Chicago Press, 2019.

Althoff, Lukas and Hugo Reichardt, "Task-specific technical change and comparative advantage," 2026. Working paper.

Appel, Ruth, Peter McCrory, Alex Tamkin, Michael Stern, Miles McCain, and Tyler Neylon, "Anthropic Economic Index report: Uneven geographic and enterprise AI adoption," Anthropic (third report, September 2025); arXiv:2511.15080 2025.

Arnon, Alexander and Kent Smetters, "Projecting the economic effects of artificial intelligence," Penn Wharton Budget Model 2025.

Autor, David and Anna Salomons, "Is automation labor-displacing?," *Brookings Papers on Economic Activity*, 2018, pp. 1–63. Spring.

Autor, David and Neil Thompson, "Expertise," *Journal of the European Economic Association*, 2025. Advance access, June 2025.

Autor, David, Caroline Chin, Anna Salomons, and Bryan Seegmiller, "New frontiers: The origins and content of new work, 1940–2018," *Quarterly Journal of Economics*, 2024, *139* (3).

Autor, David, David Dorn, and Gordon H. Hanson, "The China syndrome: Local labor market effects of import competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–2168.

Autor, David, David Dorn, Gordon H. Hanson, and Jae Song, "Trade adjustment: Worker-level evidence," *Quarterly Journal of Economics*, 2014, *129* (4), 1799–1860.

Autor, David, Frank Levy, and Richard J. Murnane, "The skill content of recent technological change: An empirical exploration," *Quarterly Journal of Economics*, 2003, *118* (4), 1279–1333.

Azar, José, Mireia Giné, and Jaime Sanz-Espín, "AI is already eroding wages: Quasi-experimental evidence from occupational exposure," SSRN working paper 5842084 2025.

Baily, Martin N., Erik Brynjolfsson, and Anton Korinek, "Machines of mind: The case for an AI-powered productivity boom," Brookings Institution, May 2023 2023.

Barattieri, Alessandro, Susanto Basu, and Peter Gottschalk, "Some Evidence on the Importance of Sticky Wages," *American Economic Journal: Macroeconomics*, 2014, *6* (1), 70–101.

Bick, Alexander, Adam Blandin, and David J Deming, "The rapid adoption of generative AI," Working Paper 32966, NBER 2024. And subsequent updates.

Blanchard, Olivier J. and Jordi Galí, "Real Wage Rigidities and the New Keynesian Model," *Journal of Money, Credit and Banking*, 2007, *39* (s1), 35–65.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, "Are ideas getting harder to find?," *American Economic Review*, 2020, *110* (4), 1104–1144.

Bonney, Kathryn, Cory Breaux, Emin Dinlersoz, Lucia Foster, John Haltiwanger, and Aakash Pande, "The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks," Working Paper CES-WP-26-25, U.S. Census Bureau, Center for Economic Studies 2026.

Braxton, J. Carter and Bledi Taska, "Technological change and the consequences of job loss," *American Economic Review*, 2023, *113* (2), 279–316.

Briggs, Joseph and Devesh Kodnani, "The potentially large effects of artificial intelligence on economic growth," Goldman Sachs Global Economics Analyst 2023.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond, "Generative AI at work," *Quarterly Journal of Economics*, 2025, *140* (2), 889–942.

Burstein, Ariel, Eduardo Morales, and Jonathan Vogel, "Changes in between-group inequality: Computers, occupations, and international trade," *American Economic Journal: Macroeconomics*, 2019, *11* (2), 348–400.

Carrillo-Tudela, Carlos and Ludo Visschers, "Supplement to "Unemployment and Endogenous Reallocation over the Business Cycle": replication files," Econometrica Supplemental Material, 91 2023.

Carrillo-Tudela, Carlos and Ludo Visschers, "Unemployment and Endogenous Reallocation over the Business Cycle," *Econometrica*, 2023, *91* (3), 1119–1153.

Caselli, Francesco and Alan Manning, "Robot arithmetic: New technology and wages," *American Economic Review: Insights*, 2019, *1* (1), 1–12.

Comin, Diego A. and Bart Hobijn, "An exploration of technology diffusion," *American Economic Review*, 2010, *100* (5), 2031–2059.

Cortes, Guido Matias, Nir Jaimovich, and Henry E. Siu, "Disappearing routine jobs: Who, how, and why?," *Journal of Monetary Economics*, 2017, *91*, 69–87.

David, Paul A., "The dynamo and the computer: An historical perspective on the modern productivity paradox," *American Economic Review*, 1990, *80* (2), 355–361.

Davidson, Tom, "What a compute-centric framework says about takeoff speeds," Open Philanthropy report, June 2023 2023.

Davidson, Tom, Basil Halperin, Trevor Houlden, and Anton Korinek, "When does automating AI research produce explosive growth? Feedback loops in innovation networks," Working paper, May 2026 2026.

Davis, Steven J. and Till von Wachter, "Recessions and the Costs of Job Loss," *Brookings Papers on Economic Activity*, 2011, *2011* (2), 1–72.

Davis, Steven J., R. Jason Faberman, and John Haltiwanger, "The establishment-level behavior of vacancies and hiring," *Quarterly Journal of Economics*, 2013, *128* (2), 581–622.

Demirer, Mert, Leon Musolff, and Luming Yang, "Writing code vs. shipping code: Productivity effects across generations of AI coding tools," Working Paper 35275, NBER May 2026.

den Haan, Wouter J., Garey Ramey, and Joel Watson, "Job Destruction and Propagation of Shocks," *American Economic Review*, 2000, *90* (3), 482–498.

Diamond, Peter A., "Wage Determination and Efficiency in Search Equilibrium," *Review of Economic Studies*, 1982, *49* (2), 217–227.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, "GPTs are GPTs: Labor market impact potential of LLMs," *Science*, 2024, *384* (6702), 1306–1308. Working paper 2023.

Erdil, Ege, Andrei Potlogea, Tamay Besiroglu, Edu Roldan, Anson Ho, Jaime Sevilla, Matthew Barnett, Matej Vrzla, and Robert Sandler, "GATE: An integrated assessment model for AI automation," arXiv:2503.04941, Epoch AI 2025.

Filippucci, Francesco, Peter Gal, and Matthias Schief, "Miracle or myth? Assessing the macroeconomic productivity gains from artificial intelligence," Artificial Intelligence Papers 29, OECD 2024.

Filippucci, Francesco, Peter Gal, and Matthias Schief, "Aggregate productivity gains from AI: A sectoral perspective," *AEA Papers and Proceedings*, 2025. Forthcoming; see also Filippucci, Gal, Laengle, and Schief (2025), Macroeconomic productivity gains from artificial intelligence in G7 economies, OECD AI Papers No. 41.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Backman, Etienne Breton, Grace Cooper, Julia A. Rivera Drew, Stephanie Richards, David Van Riper, and Kari C. W. Williams, "IPUMS CPS: Version 13.0 [dataset]," Minneapolis, MN: IPUMS 2025.

Freund, Lukas B. and Lukas F. Mann, "Job transformation, specialization, and the labor market effects of AI," 2026. Working paper.

Fukui, Masao, Emi Nakamura, and Jón Steinsson, "The commoditization of labor," 2026. Working paper, May.

Grigsby, John, Erik Hurst, and Ahu Yildirmaz, "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data," *American Economic Review*, 2021, *111* (2), 428–471.

Griliches, Zvi, "Hybrid corn: An exploration in the economics of technological change," *Econometrica*, 1957, *25* (4), 501–522.

Haefke, Christian, Marcus Sonntag, and Thijs van Rens, "Wage Rigidity and Job Creation," *Journal of Monetary Economics*, 2013, *60* (8), 887–899.

Huang, Saffron, Bryan Seethor, Esin Durmus, Kunal Handa, Miles McCain, Michael Stern, and Deep Ganguli, "How AI is transforming work at Anthropic," Anthropic, December 2025 2025.

Huckfeldt, Christopher, "Understanding the scarring effect of recessions," *American Economic Review*, 2022, *112* (4), 1273–1310.

Hulten, Charles R., "Growth accounting with intermediate inputs," *Review of Economic Studies*, 1978, *45* (3), 511–518.

Humlum, Anders, "Robot adoption and labor market dynamics," 2019. Working paper, Princeton University.

Humlum, Anders and Emilie Vestergaard, "Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI," Working Paper 33777, NBER 2025. Originally circulated as "Large Language Models, Small Labor Market Effects".

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan, "Earnings Losses of Displaced Workers," *American Economic Review*, 1993, *83* (4), 685–709.

Jones, Charles I., "R&D-based models of economic growth," *Journal of Political Economy*, 1995, *103* (4), 759–784.

Jones, Charles I. and Christopher Tonetti, "Past Automation and Future A.I.: How Weak Links Tame the Growth Explosion," Working paper, Stanford GSB, May 2026 2026.

Kambourov, Gueorgui and Iourii Manovskii, “Occupational Specificity of Human Capital,” *International Economic Review*, 2009, *50* (1), 63–115.

Karabarbounis, Loukas and Brent Neiman, “The global decline of the labor share,” *Quarterly Journal of Economics*, 2014, *129* (1), 61–103.

Kokotajlo, Daniel, Scott Alexander, Thomas Larsen, Eli Lifland, and Romeo Dean, “AI 2027,” AI Futures Project, April 2025 2025.

Korinek, Anton, “Scenario planning for an A(G)I future,” IMF Finance and Development Magazine 2023.

Korinek, Anton and Donghyun Suh, “Scenarios for the transition to AGI,” Working Paper 32255, NBER 2024.

Korinek, Anton and Joseph E. Stiglitz, “Artificial intelligence and its implications for income distribution and unemployment,” in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *The Economics of Artificial Intelligence*, NBER and University of Chicago Press, 2019, pp. 349–390.

Massenkoff, Maxim and Peter McCrory, “Labor market impacts of AI: A new measure and early evidence (“observed exposure”),” Anthropic research note, March 2026 2026.

McKinsey Global Institute, “The economic potential of generative AI: The next productivity frontier,” 2023.

Michaillat, Pascal, “Do Matching Frictions Explain Unemployment? Not in Bad Times,” *American Economic Review*, 2012, *102* (4), 1721–1750.

Moll, Benjamin, Lukasz Rachel, and Pascual Restrepo, “Uneven Growth: Automation’s Impact on Income and Wealth Inequality,” *Econometrica*, 2022, *90* (6), 2645–2683.

Mortensen, Dale T. and Christopher A. Pissarides, “Job Creation and Job Destruction in the Theory of Unemployment,” *Review of Economic Studies*, 1994, *61* (3), 397–415.

Nguyen, David, Erik Brynjolfsson, Sophia Kazinnik, Avinash Collis, and Felix Eggers, “What is generative AI worth?,” SSRN working paper 2026.

Petrongolo, Barbara and Christopher A. Pissarides, “Looking into the Black Box: A Survey of the Matching Function,” *Journal of Economic Literature*, 2001, *39* (2), 390–431.

Pilossoph, Laura, “Sectoral shocks and move unemployment,” 2012. Working paper.

Pissarides, Christopher A., *Equilibrium Unemployment Theory*, 2nd ed., Cambridge, MA: MIT Press, 2000.

Rivera-Batiz, Luis A. and Paul M. Romer, “Economic Integration and Endogenous Growth,” *Quarterly Journal of Economics*, 1991, *106* (2), 531–555.

Romer, Paul M., “Endogenous technological change,” *Journal of Political Economy*, 1990, *98* (5), S71–S102.

Tamkin, Alex and Peter McCrory, “Estimating AI productivity gains from Claude conversations,” Anthropic research report, November 2025 2025.

Trammell, Philip and Anton Korinek, “Economic growth under transformative AI,” *Annual Review of Economics*, 2026, *18*.

U.S. Bureau of Labor Statistics, “Job Openings and Labor Turnover Survey: Quits, Total Nonfarm [JTSQUR],” Retrieved from FRED, Federal Reserve Bank of St. Louis, August 19 2026.

U.S. Bureau of Labor Statistics, “Labor Force Statistics from the Current Population Survey, Household Data, Annual Averages 2025: Table 11, Employed people by detailed occupation, and Table 25b, Unemployment levels and rates by detailed occupation,” 2026.

U.S. Bureau of Labor Statistics, “Unemployment Level [UNEMPLOY] and Number Unemployed for Less Than 5 Weeks [UEMPLT5], Current Population Survey,” Retrieved from FRED, Federal Reserve Bank of St. Louis, August 19 2026.

U.S. Census Bureau, “Business Trends and Outlook Survey (BTOS): AI use estimates,” Biweekly estimates and AI supplement, releases through May 2026 2026.

Zeira, Joseph, “Workers, machines, and economic growth,” *Quarterly Journal of Economics*, 1998, *113* (4), 1091–1117.