



BIS Working Papers


No 1381

Transformative AI and demand bottlenecks

by Phurichai Rungcharoenkitkul

Monetary and Economic Department

October 2026



JEL classification: E22, E43, E52, O31, O33, O41

Keywords: artificial intelligence, automation, natural rate of interest, demand-pull innovation, market size, endogenous growth, demand bottlenecks

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in this publication are those of the authors and do not necessarily reflect the views of the BIS or its member central banks.

This publication is available on the BIS website (www.bis.org).

© *Bank for International Settlements 2026. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.*

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Transformative AI and Demand Bottlenecks*

Phurichai Rungcharoenkitkul[†]

10 October 2026

Abstract

Transformative AI promises to accelerate growth through productivity enhancement and automation. Yet automation on such a scale would raise the share of investment in output, potentially starving AI innovation of returns from a shrinking consumer market. We propose an endogenous growth framework to study this interaction, featuring research and entry decisions that propel AI advances but whose costs must be paid for by consumer demand. We identify several scenarios, spanning transformative AI, a *demand bottleneck* outcome where thinning consumer demand stalls AI progress, and intermediate cases in between. When anchored to evidence, the calibrated model points to the demand bottleneck as a highly plausible scenario, and projects more moderate growth in the coming years relative to recent studies such as Korinek et al. (2026). We draw policy implications, including for the natural rate of interest.

JEL: E22, E43, E52, O31, O33, O41.

Keywords: artificial intelligence, automation, natural rate of interest, demand-pull innovation, market size, endogenous growth, demand bottlenecks.

*I thank Claudio Borio, Denis Gorea, Pablo Hernández de Cos, Enisse Kharroubi, Giovanni Lombardo, Benoit Mojon, Daniel Rees, Matthias Rottner, Frank Smets and BIS seminar participants for helpful comments and discussions. Artificial intelligence has been used for research assistance and editorial refinement. All errors are mine. The views expressed in this publication are solely mine and do not necessarily reflect the views of the BIS or its member central banks.

[†]Bank for International Settlements. Email: phurichai.rungcharoenkitkul@bis.org

1 Introduction

Transformative artificial intelligence (AI), a type of capital capable of automating an ever wider set of tasks while continuously improving its ability, has revived the question of how fast an economy can grow. Unlike earlier technologies, AI could in principle perform virtually all cognitive tasks and improve capability shared across tasks. Some recent studies have suggested that, should AI turn transformative, annual output growth could accelerate to levels far above the historical norm of 2%, up to 15 to 18% within a few years in the most optimistic scenarios (Korinek and Suh, 2024; Korinek et al., 2026; Jones and Tonetti, 2026; Trammell and Korinek, 2026). These numbers are much larger than the more conservative estimates of 0.1-0.6 percentage points per year by Acemoglu (2024) and Filippucci et al. (2024) that rule out such transformation, at least in the short term, by assuming a fixed set of tasks AI can perform and no change in the process of invention. Transformative AI thus has the potential to fundamentally reshape the engines of economic growth.

At the same time, automation at the scale of transformative AI will likely absorb resources. As more new tasks are automated, new machines and equipment must be built and maintained, raising investment as a share of output. The composition of demand rotates away from consumption and toward accumulating and replacing capital that carries out AI services. For innovators this rotation matters, because durable market power has historically resided in consumer products, from consumer devices to cars and medicines. A critical question is whether innovation could be sustained when rewards from doing so must be earned from an ever narrowing consumption base while its costs scale with the whole economy.

In this paper, we explore the possibility of such a *demand bottleneck* as a constraint on otherwise transformative AI. We study a task-based growth model in which innovation is endogenous and carried out by innovators who compete for temporary monopoly rents. These rents accrue on consumer products, so research is funded by consumer spending rather than total output. As automation expands to cover more tasks, it thins the consumer market and the return to innovation falls. When further innovation is no longer worthwhile, automation stalls and the growth engine slows.

We formalise this mechanism in an endogenous growth model. Households save and consume as in the Ramsey problem, tasks are performed by machines or people as in Acemoglu and Restrepo (2018), and innovation is a quality ladder under free entry as in Grossman and Helpman (1991) and Aghion and Howitt (1992). A key departure is where rents come from and how long they last. Innovators earn a temporary markup on consumer products, so the prize for innovating scales with consumer spending while the cost of research scales with the economy. The consumption rate, per unit of GDP, then enters the free-entry

condition. Free entry determines an arrival rate of innovations that rises with the consumption rate and falls to zero at a threshold. At this point, the demand bottleneck binds and AI entry stops as the consumer market becomes too thin to justify further AI innovation.

Whether the bottleneck binds depends on the longevity of AI rents and the speed with which AI automates and improves capability (AI arrival rate). We identify five scenarios, a pre-AI baseline and four AI outcomes. Under “Transformative”, long-lived rents keep entry profitable despite a thinning consumer market, and trend growth increases most rapidly as the labour task share collapses. Under “Bounded boost”, machines improve at their tasks but do not displace labour, so trend growth rises only moderately and wages increase throughout. Between these two possibilities lies an intermediate scenario, “Automation without innovation”, where automation continues to completion as the innovation engine withers. Finally, “Demand bottleneck” is one where automation and innovation both stall prematurely. Our framework provides one discipline in assessing the likelihood of these scenarios. Under a calibration anchored by the available evidence, we argue that Demand bottleneck is a highly plausible one.

Under Demand bottleneck, trend growth tracks a hump-shaped path, rising temporarily before falling. Upon AI activation, innovation surges and is further amplified by automation, lifting trend growth. As automation proceeds, the required investment crowds out the consumer market and lessens incentives for new entries to challenge incumbents and sustain innovation. As a consequence, trend growth eventually falls even below its pre-AI level. This slowdown in trend growth takes time to manifest, however, masked by a persistent increase in realised output growth on the back of capital deepening associated with the AI machine build-out.

The calibrated model provides quantitative predictions of output growth, informing the current debate about growth consequences of AI. Our projections point to an acceleration of growth to around 5.5% after two decades, under the most optimistic scenario. While already impressive, these numbers are much lower than recent estimates by Korinek et al. (2026), Jones and Tonetti (2026) and Korinek and Suh (2024), where growth is projected to reach 13 to 18% within two decades, and in some scenarios within five years. We discuss the reasons for these differences: the trigger and speed of automation, and, most consequentially, whether the incentives to keep innovating are taken as given or have to be paid for by consumers.

The framework also speaks to several other debates on AI. On the labour market, the model predicts lower wages than the no-AI path for decades in every scenario with more automation, and a significant fall in the labour task share from 60% to 15-50% within three decades, depending on scenarios. The model also yields implications for monetary policy,

particularly in terms of the natural rate of interest which conceptually can be expressed as a linear function of trend growth. We highlight the large degree of uncertainty over the scenarios as well as estimation, making it especially hard to draw a definitive conclusion about the trajectory of r -star. Finally, our framework sheds light on the broader policy agenda to enhance AI benefits while limiting risks and side effects. In our model, subsidising research can offset the growth shortfall, while redistributing income between households is ineffective at averting the demand bottleneck.

Related literature The paper contributes to four literatures. The first is the growth theory of transformative AI. Aghion et al. (2019) bring AI into the task framework of Zeira (1998) and Acemoglu and Restrepo (2018) and show that growth is capped by the tasks that remain essential and hard to automate, and Jones and Tonetti (2026) quantify these weak links. Trammell and Korinek (2026) survey the models in which AI substitutes for labour in production or in research and derive the consequences for growth and interest rates, and Jones (2026) takes stock of what AI may do to growth. Korinek and Suh (2024) and Korinek et al. (2026) simulate the transition along assumed schedules for capability and adoption, Jones (2024) weighs the growth against the existential risk, and Nordhaus (2021), Erdil and Besiroglu (2023) and Davidson et al. (2026) examine the case for an acceleration, up to explosive growth from automated research. These studies take the pace of AI progress or the resources devoted to it as given and trace its consequences. This paper adds the decision to keep pushing the frontier, and the consumer spending that pays for it.

The second is the literature on demand and innovation. That demand directs invention goes back to Schmookler (1966). Consumer demand drives industrialisation in Murphy et al. (1989), the rise of mass consumption in Matsuyama (2002), the pace of innovation in Foellmi and Zweimüller (2006) and its direction in Acemoglu and Linn (2004) and Jaravel (2019). Benigno and Fornaro (2018) show that weak demand and weak innovation can lock each other in when prices are sticky, Gries and Naudé (2021) put demand constraints into the race between man and machine, and Banerjee and Wang (2026) obtain a self-fulfilling low-demand equilibrium when payrolls support demand across sectors, so that automating and hiring become strategic complements. Our rent base follows the quality-ladders-in-consumption tradition of Grossman and Helpman (1991) and Foellmi and Zweimüller (2006) rather than the total-output base of Romer (1990) and Aghion and Howitt (1992). What is new is the source of the demand shortfall: no friction and no distributional shift, only the investment that automation requires.

The third is the debate on AI and the natural rate. Trammell and Korinek (2026) convert faster growth into higher rates through the Euler equation, and Melina and Villa (2025) obtain

the same sign through the capital demand of the AI build-out when AI capital complements labour. Chow et al. (2026), Andrews and Farboodi (2025) and Kung et al. (2026) read expectations of transformative AI off real yields and asset prices, Christensen and Rudebusch (2026) find that news about AI does not explain the recent rise in estimated r^* , and Maresca (2026) argues that artificial general intelligence could lower rates instead. On the demand side, Fornaro and Wolf (2022, 2026) find that automation can depress demand and the natural rate through nominal rigidities and the shift of income from hand-to-mouth workers to capitalists with higher saving rates, in the tradition of Mian et al. (2021), Lee (2023) and Bayraktar (2026), in which inequality among savers or differences in propensities to consume depress rates or growth. We keep the natural rate equal to the discount rate plus trend growth throughout, and what falls is trend growth. The estimates of Laubach and Williams (2003) and Holston et al. (2023) anchor the level of the rate in our quantitative work.

The fourth is the evidence on what AI does to work and productivity. Task-level measures of exposure (Eloundou et al., 2024) and of adoption (Bick et al., 2026; Humlum and Vestergaard, 2025), and field experiments on productivity (Brynjolfsson et al., 2025; Noy and Zhang, 2023; Dell’Acqua et al., 2023), feed the aggregate estimates of Acemoglu (2024), Briggs and Kodnani (2023), Filippucci et al. (2024) and Aldasoro et al. (2024), which we use to bound the AI arrival rate at activation. On the labour market, Acemoglu and Restrepo (2018, 2019, 2022), Hémous and Olsen (2022), Moll et al. (2022) and Autor et al. (2024) study displacement, the creation of new tasks and the distribution of the gains, and Acemoglu and Johnson (2023) and Brynjolfsson (2022) argue for steering AI toward complementing workers. Our labour-market results share the displacement mechanism and add its feedback on the funding of innovation.

Outline The paper proceeds as follows. Section 2 lays out the model. Section 3 derives the balanced growth path, the entry threshold and the transition. Section 4 calibrates the model and simulates the five scenarios. Section 5 places the results in the AI debate. Section 6 concludes.

2 The growth model

Time is continuous. There is no aggregate uncertainty. Prices are flexible and labour is fully employed throughout. The economy is a standard growth model, with Ramsey households, task-based production, and quality-ladder innovation under free entry. The one

novel assumption pertains to the source and longevity of innovation rents.

2.1 Technology

Output requires a unit continuum of tasks, each performed either by machines or people (Acemoglu and Restrepo, 2018). Machine capital K_t with capability A_t performs the automated tasks, a fraction ϕ_t of the continuum. At an exogenous *activation* date, $t = 0$, AI becomes potentially *transformative*, i.e. purposeful AI research can continually improve machine capability and widen the set of tasks machines can perform. Effective labour H_t performs the rest, with fixed labour supply at 1 so that H_t is the labour-augmenting productivity. Tasks are complements with unit elasticity, with cost minimisation spreading each factor evenly across its tasks. Aggregate real output then reduces to

$$Q_t = \left(\frac{A_t K_t}{\phi_t} \right)^{\phi_t} \left(\frac{H_t}{1 - \phi_t} \right)^{1 - \phi_t}. \quad (1)$$

The Cobb–Douglas exponents are the measures of tasks that each factor performs, with the denominators being the factor quantities per task, $k_t = K_t/\phi_t$ and $n_t = 1/(1 - \phi_t)$.¹ Automating a task moves it from labour to machines, raising ϕ_t .² Higher average capability A makes machines better at their tasks. Unit elasticity gives every task an equal share of factor cost, so that ϕ_t is capital’s share of factor income. The automation share and the capital share are the same parameter ϕ_t .

GDP sums real output and profits from the product market

$$Y_t = Q_t + \Pi_t. \quad (2)$$

Section 2.4 details the accounting identity. In what follows, output means real output Q_t , which equals factor income with shares ϕ_t and $1 - \phi_t$. The consumption, investment and research shares are expressed relative to GDP, Y .³ Capital is accumulated from output and depreciates at rate δ :

$$\dot{K}_t = I_t - \delta K_t. \quad (3)$$

¹Equation (1) nests two geometric means. For each product, output is the average of its task outputs, with the product’s own indices. For the economy, real output is the average of product outputs, with A_t and H_t the geometric means of the product indices. See Appendix A.1 for detailed aggregation.

²Conventional automation, including that powered by non-transformative AI, continues after activation but is offset by new tasks for labour, resulting in zero net displacement and a constant automated share, ϕ_0 , which we interpret as the exhausted frontier of automation before transformative AI.

³The labour share is the task share $1 - \phi_t$. Labour’s share of GDP is smaller at $(1 - \phi_t)(1 - \mu/\varepsilon)$, with μ the consumption rate and ε the demand elasticity of Section 2.4.

On a balanced path the automated share is constant at ϕ and capital grows with output. The growth rate under (1) is then given by

$$g = \frac{\phi}{1-\phi} g_A + g_H, \quad (4)$$

where g_A and g_H are the growth rates of the two productivity indices. AI affects trend growth through its influence on (ϕ, g_A, g_H) .

2.2 Automation dynamics

Automation ϕ_t follows capability A_t with a diffusion lag, separating the invention of automation from its rollout. The automatable frontier ϕ_t^* is the share of tasks machines could perform at current capability. We assume capability growth $g_{A,t}$ advances the frontier in proportion to its remaining gap to the ceiling, while adoption closes the gap between the automated share ϕ_t and the frontier at rate ξ :

$$\dot{\phi}_t^* = \psi (\phi_{\max} - \phi_t^*) g_{A,t}, \quad \phi_{\max} < 1, \quad (5)$$

$$\dot{\phi}_t = \xi (\phi_t^* - \phi_t). \quad (6)$$

The ceiling $\phi_{\max}$ is an assumed supply-side limit to automation, in the spirit of Aghion et al. (2019) and Jones and Tonetti (2026). The adoption lag reflects the historical tendency of technologies to reach wide use long after they become feasible (Comin and Hobijn, 2010). Appendix A.2 shows that the rule stands in for the choice competitive producers would make, where they keep automating after AI innovation stops and never switch back. Before the activation date, $g_A = 0$ and $\phi^* = \phi = \phi_0$, so the economy is on its pre-AI balanced path, where H drives all progress.

The formulation subsumes different AI scenarios after activation. AI may be used as a tool to raise worker productivity H at a non-automated task with $g_A = 0$ and $g = g_H > 0$, while leaving the automated share unchanged. In this paper, we regard such uses as part of conventional innovation (see Section 2.5). As with earlier technologies, from the industrial era to recent robotics, automation coincides with the creation of new labour tasks, keeping the factor shares stable. Transformative AI would instead sustain capability improvement through continuing research, so that $g_A > 0$ and the frontier keeps advancing, $\dot{\phi}_t^* > 0$ in (5), as automation outpaces the arrival of new labour tasks. How long the automation process continues is an outcome of free entry into AI research, specified in Section 2.5. Our definition of transformative AI does not include automated research, as capability grows only with new

entry and capability does not itself make research cheaper. In this respect it overlaps with, but is narrower than, the notion surveyed in Trammell and Korinek (2026), which includes research automation. Appendix C.2 considers the implications of automated research within our setup.

2.3 Households

Two types of households populate the economy. A unit measure of *capitalists* owns all capital and firm equity, with full access to asset markets. They choose consumption to maximise

$$U_k = \int_0^\infty e^{-\rho t} \frac{C_{k,t}^{1-\sigma} - 1}{1-\sigma} dt, \quad \rho > 0, \sigma > 0, \quad (7)$$

subject to the flow budget constraint

$$\dot{W}_t = r_t W_t - C_{k,t}, \quad (8)$$

where W is financial wealth (capital plus firm equity) and r_t is the common required return delivered by no-arbitrage across the assets they hold, together with the transversality condition $\lim_{t \rightarrow \infty} e^{-\int_0^t r_s ds} W_t = 0$. *Workers* supply one unit of labour inelastically, have no access to asset markets, and consume their income each period, $C_{w,t} = (1 - \phi_t)Q_t$.

The first-order condition for the capitalist's problem gives the standard Euler equation

$$\frac{\dot{C}_{k,t}}{C_{k,t}} = \frac{r_t - \rho}{\sigma}. \quad (9)$$

In what follows, we specialise to log utility, $\sigma = 1$, under which the capitalist consumes the annuity value of wealth, $C_k = \rho W$. We keep general σ where a comparison with the literature is helpful, as in the balanced-path rate of return, $r = \rho + \sigma g$, which anchors the natural rate of interest in Section 4.

Financial wealth consists of capital and firm equity,

$$W_t = K_t + V_t \quad (10)$$

where the equity claims V_t is valued in the aggregate, earning aggregate profits Π_t net of the research outlay R_t that the same households fund:

$$\dot{V}_t = r_t V_t - (\Pi_t - R_t) \quad (11)$$

with Π_t and R_t described below.

2.4 Rents

Households allocate consumer spending $C_t = C_{k,t} + C_{w,t}$ over a unit continuum of differentiated products, with unit elasticity of substitution across products. Within its own market, each product faces an isoelastic demand with elasticity $\varepsilon > 1$. The latest innovator on each product is its monopolist. It makes its product with the task technology of Section 2.1, automated to the same extent as every other product but with its own machine capability and labour productivity, and sets prices at the constant markup $\varepsilon/(\varepsilon - 1)$ over its unit cost. Equal spending at a common markup implies the same cost outlay for every product, and hence the same capital and labour. A product with a lower unit cost sells more units at a lower price from the same inputs. Equation (1) therefore holds in aggregate with A_t and H_t being the geometric means across products (Appendix A.1).

Of every dollar spent on a product, $(\varepsilon - 1)/\varepsilon$ covers the production cost and $1/\varepsilon$ accrues to the incumbent as profit. Aggregate profits over all products are

$$\Pi_t = \frac{1}{\varepsilon} C_t = \frac{1}{\varepsilon} \mu_t Y_t, \quad (12)$$

where $\mu_t \equiv C_t/Y_t$ is the consumption rate, namely total consumption per unit of GDP. With a unit measure of products, each incumbent's profit flow π_t equals the aggregate:

$$\pi_t = \Pi_t = \frac{1}{\varepsilon} \mu_t Y_t. \quad (13)$$

An incumbent earns the flow π_t while its monopoly lasts, whereas the sector earns Π_t permanently.

Capital goods are produced competitively and carry no rent. Competitive task pricing awards each factor its marginal product of real output,

$$r_t = \phi_t \frac{Q_t}{K_t} - \delta, \quad (14)$$

$$w_t = (1 - \phi_t) \frac{Q_t}{H_t}, \quad (15)$$

with $Q_t = Y_t - \Pi_t$.

GDP exceeds real output by exactly the retail margin on consumption. The economy produces a single physical flow Q and allocates it to consumption, investment and research at producer prices. Investment and research goods, I and R , are the same composite of products,

sold at cost. Consumers spend C on products, Π of which goes to firms as a transfer due to a markup embedded in the retail price. The same allocation, valued at the two price vectors, produces a goods market-clearing condition and the GDP identity:

$$Q_t = (C_t - \Pi_t) + I_t + R_t, \quad \text{real output at cost,} \quad (16)$$

$$Y_t = C_t + I_t + R_t, \quad \text{GDP at market prices.} \quad (17)$$

On the income side, wages $(1 - \phi)Q$, capital rentals ϕQ and profits Π sum to the same $Q + \Pi$. This distinction between output and GDP is standard national accounting.⁴ On any balanced path μ is constant, so Q and Y grow at the same rate, and every growth and natural-rate statement is independent of which aggregate one uses. The distinction nonetheless matters for growth, because markups incentivise innovation. As the markup falls on consumption alone, the reward to innovation (profit flow (13)) is proportional to consumer spending μY rather than total output Y . If the rent base were instead total output, the profit flow would be Y/ε independent of the spending composition. No demand bottleneck could arise, as automation could never thin the market relative to the economy. So while the difference between Y and Q does not affect growth measurement, it is decisive for whether growth continues.

A key assumption is that innovation rents originate in the consumer market.⁵ Enduring franchises are historically tied to consumer products, such as patented medicines, branded goods, and consumer devices together with associated services. By contrast, capital such as machine tools, memory chips or fibre optics tends to be competitively supplied, as evidenced by the falling relative price of equipment consistent with the loss of pricing power (Greenwood et al., 1997). The internet was a recent example, with durable rents accruing to consumer-facing businesses rather than the network or equipment sellers. We carry this assumption over to the current AI episode. Appendix A.3 considers the choice between the two more formally.

⁴A markup is a private excise tax on consumer goods whose revenue accrues to the incumbents, so Q is GDP at factor cost and Y is GDP at market prices. The revenue-to-cost wedge is the object the markup-measurement literature estimates (Hall, 1988; Baqaee and Farhi, 2020), and the model's profit share $\Pi/Y = \mu/\varepsilon$ is its counterpart.

⁵This reverses the canonical assumption, in which the rent base scales with total output, as in Romer (1990) or Aghion and Howitt (1992). Our consumer-side rent base follows the quality-ladders-in-consumption tradition (Grossman and Helpman, 1991; Foellmi and Zweimüller, 2006), and market-size arguments in which the consumer market determines industrialisation (Murphy et al., 1989), the rise of mass consumption (Matsuyama, 2002), and the direction of innovation (Jaravel, 2019).

2.5 Innovation

Monopoly rents are temporary due to two forces. First, rents are subject to exogenous obsolescence at rate $d_x > 0$, due to imitation, patent expiry or preference turnover. Second, monopoly power is exposed to creative destruction at rate ι_x from entry of the next innovator. The product itself remains supplied as in Section 2.4 and aggregate profits stay at C/ε , but the franchise value transfers among the capitalists.⁶ The creative destruction hazard ι generates endogenous growth, while the obsolescence hazard d is what will stall growth (Section 3.2).

Innovation is directed research on two production processes, indexed $x \in \{A, H\}$. AI research advances the product's machine capability A_j while conventional research improves its labour productivity H_j . To the extent that AI also improves labour productivity, we treat it as part of the conventional innovation, no different from earlier tools (we relax this assumption in Section 5.2). Innovators engage in research to win monopoly power over one of these two processes. The product's rent is then the sum of the two rents corresponding to the machine and labour processes, earned separately and with fixed shares $\alpha_A + \alpha_H = 1$. The entry problem takes the same form along both dimensions. An entrant pays a flow cost $\alpha_x \eta_x Y_t$ per unit arrival rate of innovations.⁷ A success raises the productivity of the process it improves by the step θ_x and transfers the rent on that process to the entrant, displacing the incumbent.⁸ Each process is thus exposed to its own dimension's creative destruction ι_x only.⁹ The machine rent held before activation is not contested, as the scope for improving machine processes is exhausted.

The ability to conduct AI research spreads gradually, product by product. We assume the share of products whose machine process can benefit from purposeful AI research, s_t , expands with a lag:

$$\dot{s}_t = \xi (1 - s_t), \quad s_0 = 0, \quad (18)$$

⁶Letting obsolete products lose their markup instead of passing to a new monopolist would not change the prize for innovation, which is a fixed share of consumer spending on the innovator's product. Under unit elasticity that spending does not depend on what other products cost, so total profits would fall but the prize for a new innovation would not. If consumers instead shifted spending towards the cheaper products, the prize would shrink as innovation slows, which reinforces the bottleneck.

⁷Hence research costs scale with the economy, which removes the scale effect identified by Jones (1995), as in Young (1998) and Howitt (1999). Appendix C.1 examines the alternative denomination of research costs in researcher wages. Aggregate research outlay is broad innovation investment in the intangibles sense (Corrado et al., 2009), of which measured R&D is the subset.

⁸The innovator holds the exclusive right to the process it improves, so the displaced holder cannot challenge it and the markup does not depend on the step size.

⁹We do not model firm boundaries. The holder of a process rent may be a supplier with a unique technology, a licence holder, or an innovator whose firm has been acquired by the product's owner. The model is consistent with all these interpretations.

while the labour process is open to conventional research at all times. The lag stands for the time it takes to build the skills, compute and organisation that AI research on a product requires, assumed to be the same as the adoption speed ξ . On any balanced path $s = 1$. As successes arrive at rate ι_x per product open to research, the two productivity indices grow at rates:

$$g_A = \theta_A s_t \iota_A, \quad (19)$$

$$g_H = \theta_H \iota_H. \quad (20)$$

Sustaining arrivals at rate $\iota_{x,t}$ costs $\alpha_x \eta_x \iota_{x,t} Y_t$ per product, so aggregate research spending is

$$R_t = \alpha_A \eta_A s_t \iota_{A,t} Y_t + \alpha_H \eta_H \iota_{H,t} Y_t, \quad (21)$$

with the arrival rates on open products determined by free entry below. The share s_t thus enters in two places: (i) it scales capability growth, pushing out the frontier in (5), and (ii) it scales the research outlay, thinning the consumption rate.

Entry is unrestricted and continues until the prize no longer covers the cost. Let $V_{x,t}$ denote the value of the rent on process x of a product, namely the present value of its share $\alpha_x \pi_t$ of the profit flow (13), discounted at r_t , until the claim ends at the hazard $d_x + \iota_{x,t}$. Meanwhile, aggregate equity V_t is discounted at r_t with no terminal hazard, as obsolescence and creative destruction merely transfer claims among capitalists. This value satisfies the no-arbitrage pricing condition:

$$r_t V_{x,t} = \alpha_x \pi_t + \dot{V}_{x,t} - (d_x + \iota_{x,t}) V_{x,t}. \quad (22)$$

Free entry then implies

$$V_{x,t} \leq \alpha_x \eta_x Y_t, \quad (23)$$

with equality whenever $\iota_{x,t} > 0$. The only reward to research, on either dimension, is a share of the consumer-product rent of (13). This assumption extends the demand-pull view of innovation (Schmookler, 1966; Acemoglu and Linn, 2004) to all of trend growth, including the pre-AI period.

An equilibrium, given (K_0, A_0, H_0, ϕ_0) with $\phi_0^* = \phi_0$ and $s_0 = 0$, is a path on which households optimise ((9) with transversality), factors are priced by (14), the values of process rents obey (22) with entry (23), automation follows (5)–(6), and research access (18), and goods clear by (17) with (3). In Section 3.3, we solve the path as a temporary equilibrium with static expectations, in which households and entrants optimise at each date and markets

clear under the belief that current conditions persist. Thus entrants price process rents off the balanced path at the current automated share ϕ_t and the current share of products open to AI, s_t , and the consumption rate takes its balanced-path value at those shares. Appendix B.1 solves the perfect-foresight equilibrium.

3 Growth and bottlenecks

3.1 The balanced growth path

On a balanced growth path the automated share is constant at ϕ , every aggregate grows at a common rate g , and the consumption rate μ is constant. We solve for the pair (μ, g) in four steps. The household problem pins the equilibrium interest rate, mapping growth into a required return. The return prices capital and incumbencies, expressing investment and the innovation prize in terms of μ and g . Free entry and goods-market clearing conditions then provide two restrictions that determine μ and g , identifying the balanced path. Automation shifts the parameter ϕ , which we come back to in Section 3.2.

Households choose consumption and pin the interest rate. On the balanced path consumption grows at g , so the Euler equation (9) implies a constant required return,

$$r = \rho + \sigma g. \quad (24)$$

Under log utility the plan $C_k = \rho W$ of Section 2.3 is exact, since it implies $\dot{W} = (r - \rho)W$ and transversality rules out every other initial level. The capitalist consumes the annuity value of wealth, whose components are priced by the same interest rate. The equity component of wealth follows from its flow valuation in Section 2.3, with $\dot{V} = gV$ on the balanced path, so

$$V = \frac{\Pi - R}{r - g}. \quad (25)$$

This aggregate is not the sum of incumbency values. Creative destruction claims are bought at full value, while obsolescence claims are free windfalls. Equation (25) exceeds summed incumbency values by the capitalised value of those windfalls.

The equilibrium interest rate prices capital. The marginal product condition (14) and factor income $Q = (1 - \mu/\varepsilon)Y$ from (2) and (12) imply a capital-output ratio $K/Y = \phi(1 - \mu/\varepsilon)/(r + \delta)$. Capital grows at g , so gross investment is $(g + \delta)K$ and the investment

share is

$$\frac{I}{Y} = (1 - a) \phi \left(1 - \frac{\mu}{\varepsilon}\right), \quad (26)$$

where

$$a \equiv \frac{r - g}{r + \delta} = \frac{\rho}{\rho + g + \delta} \quad (27)$$

is the fraction of capital's gross return left over after accumulation gK and replacement δK , where the second equality uses log utility. Automation raises the investment share directly, as a larger ϕ puts more of production on the factor that must be accumulated and replaced.

The same interest rate prices rents, and free entry determines the arrival rate of innovations. The rent on process x is a claim on the share α_x of the product rent $\pi = \mu Y / \varepsilon$ for as long as the claim survives. On the balanced path, this rent grows at g . The value that capitalises the rent grows at the same rate, so $\dot{V}_x = g V_x$. Substituting this into (22) gives

$$V_x = \frac{\alpha_x \pi}{r - g + d_x + \iota_x} = \frac{\alpha_x \mu Y / \varepsilon}{\rho + d_x + \iota_x}, \quad (28)$$

where the second step uses $r = \rho + g$ under log utility. This value is the prize an entrant wins by displacing the incumbent. More entry raises the hazard ι_x and lowers the prize, so arrivals settle where the prize just covers the entry cost, $V_x = \alpha_x \eta_x Y$. Substituting the rent (13) and cancelling Y :

$$\iota_x(\mu) = \max \left\{ 0, \frac{\mu}{\varepsilon \eta_x} - (\rho + d_x) \right\}. \quad (29)$$

The arrival rate rises with the consumption rate and falls to zero at the threshold:

$$\hat{\mu}_x \equiv \varepsilon \eta_x (\rho + d_x). \quad (30)$$

Below the threshold the prize falls short of the cost even with no competing entry, prompting innovation on that dimension to stop.

The presence of the consumption rate μ in (29) is the model's central departure from standard growth theory. It survives the cancellation of Y because the prize and the cost scale with different aggregates. The entry cost scales with the economy, $\alpha_x \eta_x Y$, while the prize scales with consumer spending through the rent $\alpha_x \mu Y / \varepsilon$. Had rents scaled with total output instead, as in Romer (1990) or Aghion and Howitt (1992), the prize of innovation would have grown in step with its cost, leaving the arrival rate constant at $1/(\varepsilon \eta) - (\rho + d)$, and insulating growth from the composition of spending. No demand bottleneck could thus arise in the standard growth model.

Goods-market clearing closes the system. Expressing (17) as a share of GDP gives $\mu = 1 - I/Y - s_R$ where $s_R \equiv R/Y$ is the research outlay share, with $s_R = \alpha_A \eta_A \iota_A + \alpha_H \eta_H \iota_H$ from (21) at $s = 1$.¹⁰ Substituting the investment share (26) and solving for μ gives the first relation, while substituting the arrival rates (29) into the growth skeleton (4) gives the second:

$$\mu = \frac{1 - (1 - a)\phi - s_R}{1 - (1 - a)\phi/\varepsilon}, \quad \text{the resource relation,} \quad (31)$$

$$g = \frac{\phi}{1 - \phi} \theta_A \iota_A(\mu) + \theta_H \iota_H(\mu), \quad \text{the entry relation.} \quad (32)$$

The balanced path is thus the intersection of the two relations. With μ declining with g from (31) and g increasing in μ from (32), a unique fixed point is guaranteed. The resource relation is smooth. The entry relation is piecewise linear, because of the shut-off in (29). Each dimension adds to growth only while μ exceeds its threshold $\hat{\mu}_x$, inducing a kink in the relation.

Automation moves both relations, pulling growth in opposite directions. A higher ϕ shifts the resource relation left, holding trend growth and the research share fixed,

$$\frac{\partial \mu}{\partial \phi} = - \frac{(1 - a) [(\varepsilon - 1) + s_R]/\varepsilon}{(1 - (1 - a)\phi/\varepsilon)^2} < 0, \quad (33)$$

as capital deepening lowers consumption at every growth rate. At the same time, a higher ϕ steepens the AI segment of the entry relation through the amplifier $M(\phi) \equiv \phi/(1 - \phi)$ in (32). The amplifier captures the two ways capability gains raise output: directly on the automated tasks, with weight ϕ , and through the capital accumulation that higher output induces, which compounds the direct effect by $1/(1 - \phi)$. It is convex, with $M'(\phi) = 1/(1 - \phi)^2 > 0$ and $M(\phi) \rightarrow \infty$ as $\phi \rightarrow 1$. This is what makes automation transformative, as the more tasks machines perform, the more a given capability advance raises growth.

A tension arises from the effects of the automated share ϕ on both the amplifier and the investment share (26). On the one hand, each task automated by machines makes capability advances add more to growth. On the other hand, automation commits more output to capital accumulation and replacement, bidding real resources away from the consumer market that sustains entry in (29). A technology advanced enough to exploit the amplifier must also fund the capital deepening that thins the consumer market. Which force prevails depends

¹⁰Equation (31) is implicit in μ , since s_R depends on μ through (29). On each branch of (29) the research share is affine in μ , so the relation solves in closed form branch by branch, and because the research share is non-decreasing in μ the solution is unique.

on where the consumption rate stands. When μ exceeds the AI threshold $\hat{\mu}_A$, a steeper AI segment competes against the resource relation moving left. If μ lies below the threshold, AI entry stops and the AI term of (32) vanishes, with no further growth gain. Past this point, automation only deepens capital and thins the consumer market further, prompting growth to fall.

Figure 1 shows how the entry and resource relations jointly determine the equilibrium for given automation shares. The pre-AI relations (i.e. before transformative AI) are represented by dashed lines, which cross at the initial balanced path of 2% growth. At this initial point, the automated share rests at ϕ_0 , and all productivity progress, including from AI being used as a complementary tool, occurs through H . Accordingly, the entry relation (32) contains only the conventional term $\theta_H \iota_H$, which rises linearly with μ for $\mu > \hat{\mu}_H$ under (29). This inequality must hold under positive balanced-path growth, with $\hat{\mu}_H$ kept low by slowly expiring conventional rents (small d_H), whose income repays entry even from a thin market. As such, the entry schedule exhibits only one kink, at $\mu = \hat{\mu}_H$. The resource relation (31) contains no AI research outlays.


The activation of transformative AI has non-monotone effects on the balanced growth path, varying with the extent of automation. Transformative AI adds the capability term to the entry relation (32), steepening the schedule where the consumption rate μ exceeds the threshold $\hat{\mu}_A$ (the right dotted vertical line) and introducing a second kink. Meanwhile, AI research outlays shift the resource relation leftward. On impact, the standing consumption rate already exceeds $\hat{\mu}_A$, so AI entry immediately pays and balanced-path growth jumps, accompanied by a slight drop in μ . But as the automated share ϕ increases, the leftward shift in the resource relation dominates the steepening of the entry relation, prompting a decline in balanced-path growth with a more notable drop in μ . The slide stops when μ reaches $\hat{\mu}_A$ and AI entry no longer pays, i.e. the demand bottleneck binds. The frontier then freezes as $g_A = \theta_A \iota_A = 0$, and the share comes to rest. Growth is then $\theta_H \iota_H$, below its starting point.

3.2 Steady-state automation

A steady state after transformative AI activates is a balanced growth path with a constant automated share. Imposing $\dot{\phi}^* = \dot{\phi} = 0$ in (5)–(6) yields $\phi = \phi^*$, so that either the frontier is saturated, $\phi = \phi_{\max}$, or AI capability growth g_A is zero. By $g_A = \theta_A \iota_A$, the latter case implies $\iota_A(\mu(\phi)) = 0$ and, by (29)–(30), means $\mu(\phi) \leq \hat{\mu}_A$. Because $\mu'(\phi) < 0$ by (33), this second set, when non-empty, is an interval bounded from below by the unique $\bar{\phi}$ that solves

$$\mu(\bar{\phi}) = \hat{\mu}_A. \quad (34)$$

Figure 1: Balanced-path growth determination



Note: The figure plots balanced-path growth g against the consumption rate μ , under an illustrative calibration. Blue curves are the entry relation (32), kinked at the thresholds $\hat{\mu}_A$ and $\hat{\mu}_H$, with the AI segment steepening as the automated share ϕ rises. Red curves are the resource relation (31), shifting left as capital deepening reduces consumption. Dashed lines are the pre-AI relations, the entry relation without an AI segment and the resource relation without AI research outlays. Solid lines are the post-activation relations, at different automated shares. Dotted verticals mark the thresholds.

When $\phi_0 < \bar{\phi} < \phi_{\max}$, there is a continuum of long-run steady states $[\bar{\phi}, \phi_{\max}]$. When $\bar{\phi} \geq \phi_{\max}$, the interval is empty and the long-run share is unique at the ceiling, $\phi = \phi_{\max}$. When $\bar{\phi} \leq \phi_0$, the consumption rate is already at or below the threshold at activation, so AI entry never starts and the automation share stays at ϕ_0 .

We solve for $\bar{\phi}$ in closed-form. Evaluating the arrival rates (29) at $\mu = \hat{\mu}_A$,

$$\iota_A(\hat{\mu}_A) = 0, \quad (35)$$

$$\iota_H(\hat{\mu}_A) = \frac{\hat{\mu}_A}{\varepsilon\eta_H} - (\rho + d_H) = \frac{\hat{\mu}_A - \hat{\mu}_H}{\varepsilon\eta_H} > 0. \quad (36)$$

The first evaluation follows from the definition of $\hat{\mu}_A$, which implies exactly zero AI entry at this crossing point. The second evaluation uses the same definition for the conventional threshold, positive under the assumption that conventional incumbencies expire more slowly,

$\hat{\mu}_H < \hat{\mu}_A$, maintained throughout.¹¹ Substituting the two arrival rates into the entry relation (32) and the research outlay (21) gives growth and the research share,

$$\hat{g} = \frac{\phi}{1-\phi} \theta_A \cdot 0 + \theta_H \iota_H(\hat{\mu}_A) = \theta_H \frac{\hat{\mu}_A - \hat{\mu}_H}{\varepsilon \eta_H}, \quad (37)$$

$$\hat{s}_R = \alpha_H \eta_H \iota_H(\hat{\mu}_A) = \frac{\alpha_H(\hat{\mu}_A - \hat{\mu}_H)}{\varepsilon}. \quad (38)$$

Finally, substituting $\mu = \hat{\mu}_A$, $s_R = \hat{s}_R$ and $g = \hat{g}$ into the resource relation (31) yields $\bar{\phi}$ in closed form:

$$\bar{\phi} = \frac{1 - \hat{\mu}_A - \alpha_H(\hat{\mu}_A - \hat{\mu}_H)/\varepsilon}{(1 - \hat{a})(1 - \hat{\mu}_A/\varepsilon)}, \quad (39)$$

where $\hat{a} = \rho/(\rho + \hat{g} + \delta)$ and $\hat{g} = \theta_H(\hat{\mu}_A - \hat{\mu}_H)/(\varepsilon \eta_H)$.

A lower $\bar{\phi}$ implies more severe demand bottlenecks, as thinning demand stalls automation earlier. The closed-form $\bar{\phi}$ reveals the conditions under which this is the case. Taking partial derivatives of (39) yields:

$$\begin{aligned} \frac{\partial \bar{\phi}}{\partial \hat{\mu}_A} &= - \frac{1 + \alpha_H/\varepsilon - (1 - \hat{a})\bar{\phi}/\varepsilon}{(1 - \hat{a})(1 - \hat{\mu}_A/\varepsilon)} < 0, \\ \frac{\partial \bar{\phi}}{\partial \hat{\mu}_H} &= \frac{\alpha_H/\varepsilon}{(1 - \hat{a})(1 - \hat{\mu}_A/\varepsilon)} > 0, \quad \frac{\partial \bar{\phi}}{\partial \hat{a}} = \frac{\bar{\phi}}{1 - \hat{a}} > 0. \end{aligned} \quad (40)$$

To gain intuition, re-state (39) as $1 = \hat{\mu}_A + \hat{s}_R + (1 - \hat{a})\bar{\phi}(1 - \hat{\mu}_A/\varepsilon)$, i.e. the output share is divided into consumption $\hat{\mu}_A$, research $\hat{s}_R$, and investment (the investment share (26) at the stall). The common denominator $(1 - \hat{a})(1 - \hat{\mu}_A/\varepsilon)$ of the first two derivatives is the investment share generated by one unit of automated tasks. In the first derivative, a higher $\hat{\mu}_A$ raises the consumption share one for one and the research share by α_H/ε , and lowers the existing investment share by $(1 - \hat{a})\bar{\phi}/\varepsilon$. The numerator is positive because investment over real output lies below one while ε exceeds one. In the second derivative, a higher $\hat{\mu}_H$ affects only the research share, lowering it by α_H/ε and freeing output to accommodate a larger automated share. In the third, a higher $\hat{a}$ means every automated task ties up less output in investment, freeing room to accommodate a larger automated share.

To trace the effects of primitive parameters on $\bar{\phi}$ through the thresholds and the ratio $\hat{a}$ in (40), we apply the chain rule. A higher d_A or η_A raises $\hat{\mu}_A$ by (30) and, by widening the threshold gap in (37), raises $\hat{g}$ and so lowers $\hat{a}$. Both channels thus lower $\bar{\phi}$. Meanwhile, a

¹¹This is consistent with empirical evidence as discussed below.

higher d_H or η_H increases $\hat{\mu}_H$ and, by narrowing the gap, raises $\hat{a}$ and hence $\bar{\phi}$. A higher θ_H or δ lowers $\hat{a}$ and so lowers $\bar{\phi}$, while ρ moves the thresholds and the ratio in opposite directions, producing an ambiguous sign. The demand bottleneck is therefore more severe when AI rents expire faster and AI entry is costlier, raising the required consumer base to maintain AI research, and when conventional research runs larger and capital wears out faster, claiming more output for each automated task.

Figure 2(a) illustrates the comparative statics. The consumer market curve $\mu(\phi)$ falls as automation proceeds, crossing the AI threshold $\hat{\mu}_A$ at the floor $\bar{\phi}$. The interval between the floor and the ceiling defines the stalled steady states, within which the consumption rate no longer repays AI entry. The interval's width then indicates the severity of the bottleneck problem, and is shaped by the forces in (40). Faster-expiring AI rents and costlier AI entry lift the threshold, larger conventional research and faster capital wear lower the market curve, and either shift moves the crossing left and widens the interval. When instead the threshold lies below the market curve at the ceiling, the curves never cross, the interval is empty, and automation can come to a rest only at $\phi_{\max}$.

3.3 Transition


The system converges to a unique steady state from any initial condition with no indeterminacy. Given the automated share at activation and the parameters, the equilibrium path is unique, and both automation shares rise monotonically, converging to a unique point in the steady-state set $[\bar{\phi}, \phi_{\max}]$. The reason is that the dynamical system contains no jump variables. At each date, the resource relation (31) sets the consumption rate given ϕ_t and s_t , free entry (29) turns it into an arrival rate on those products, the arrival rate advances the frontier, adoption narrows the gap, and access (18) unlocks further products. From any starting point there is a unique path to a unique steady state.

The transition system is the automation block (5)–(6) and the access equation (18), combined with equilibrium entry at the current state. This quasi-static specification lets the consumption rate and the arrival rate take their balanced-path values at current ϕ_t and s_t :

$$\mu_t = \mu(\phi_t, s_t), \quad \iota_{A,t} = \iota_A(\mu_t), \quad g_{A,t} = \theta_A s_t \iota_{A,t}, \quad (41)$$

where $\mu(\phi, s)$ solves the resource relation (31) with the research share $s_R = \alpha_A \eta_A s \iota_A + \alpha_H \eta_H \iota_H$ and the entry relation (32). The specification is exact at every steady state, and the forward-looking alternative is taken up in Section 4.2. The system (5)–(6) with (18) is then autonomous in (ϕ^*, ϕ, s) from the initial state $\phi^* = \phi = \phi_0$ and $s = 0$, and is depicted in Figure 2(b).

Figure 2: Steady states and convergence



Note: Illustrative schematics of steady-state determination and convergence paths. Panel (a) plots the balanced-path consumption rate $\mu(\phi)$ against the automated share and the AI threshold $\hat{\mu}_A$. Their crossing is the floor $\bar{\phi}$, and the shaded band is the steady-state interval $[\bar{\phi}, \phi_{\max}]$. Dashed schedules consider alternative parameters: the threshold under a higher d_A or η_A and the red curve under a higher θ_H or δ or a lower d_H or η_H . Open circles mark the new crossing points. Panel (b) plots the automatable frontier ϕ^* against the automated share ϕ along the automation dynamics (5)–(6), with AI entry at the current consumption rate and the steady states on the diagonal in red. All paths start at activation, and the speed of invention relative to adoption, $\psi\theta_A/\xi$, selects the case. Under slow invention (light blue) the frontier stays close to the installed share and entry stops at the edge $\bar{\phi}$. Under fast invention (dark blue) the frontier runs ahead, entry stops when the installed share reaches $\bar{\phi}$ while the frontier is already inside the interval, and adoption pushes the share to an interior steady state. Under permanent rents (green), $d_A = 0$, the interval is empty and automation completes at the ceiling.

Dividing the frontier equation by the adoption equation, wherever $\phi^* > \phi$,

$$\frac{d\phi^*}{d\phi} = \frac{\psi\theta_A}{\xi} \frac{(\phi_{\max} - \phi^*) s_t \iota_A(\mu(\phi_t, s_t))}{\phi^* - \phi}. \quad (42)$$

The trajectory and the terminal point are then functions of only ϕ_0 and $\psi\theta_A/\xi$.

The transition path exists and is unique. The consumer market curve $\mu(\phi, s)$ is continuous and piecewise smooth in both arguments, ι_A is the larger of zero and an affine function of the consumption rate, and the share of products open to AI research s_t follows a linear equation. Thus the system is Lipschitz continuous on $[\phi_0, \phi_{\max}]^2 \times [0, 1]$, and the Cauchy–Lipschitz theorem guarantees exactly one solution from the given state.¹² The path is monotone, bounded, and therefore convergent to a steady state. The share s_t rises to one, and the

¹²Determinacy here is more immediate than in the neoclassical growth model, where the consumption jump is pinned by saddle-path stability. Our system has no jump variable and expectations play no role.

frontier always weakly expands, $\dot{\phi}^* \geq 0$. The shares start equal and $\dot{\phi} = 0$ on the diagonal, so $\phi \leq \phi^*$ and hence $\dot{\phi} \geq 0$ throughout. As both shares rise and are capped by $\phi_{\max}$, they must converge. At the limit, where all motion ceases, $s = 1$ and $\phi_\infty^* = \phi_\infty$ with $\phi_\infty = \phi_{\max}$ or $\iota_A(\mu(\phi_\infty, 1)) = 0$, a steady state of Section 3.2.

The trajectory is uniquely pinned by the initial activation share and parameters. Figure 2(b) shows three types of trajectories, all starting from the same activation share ϕ_0 and differing in the parameters, which select which case obtains. If $\bar{\phi} \geq \phi_{\max}$, as under permanent rents ($d_A = 0$), then $\iota_A(\mu(\phi, s)) > 0$ at every share below the ceiling, so the only steady state is the ceiling $\phi_\infty = \phi_{\max}$. If $\bar{\phi} < \phi_{\max}$ and the path has a date at which $\phi_t^* > \bar{\phi} \geq \phi_t$, the frontier passes the floor first. The installed share then reaches the floor at a finite date T , because adoption (6) runs at a rate bounded below, $\dot{\phi} \geq \xi(\phi^* - \bar{\phi}) > 0$. At T the consumption rate reaches the threshold, $\mu(\phi_T, s_T) = \hat{\mu}_A$, entry stops, $\iota_A = 0$, and the frontier freezes at $\bar{\phi}^* \equiv \phi_T^* > \bar{\phi}$. The floor is the same at every s_t , because AI entry is zero at the threshold and s_t enters (41) only through AI entry. Adoption closes the remaining gap, $\phi_t \rightarrow \bar{\phi}^*$, the consumption rate falls strictly below the threshold, $\mu(\phi_t, s_t) < \hat{\mu}_A$, and entry never restarts. Hence, $\phi_\infty = \bar{\phi}^*$, an interior steady state. If instead $\phi_t^* \leq \bar{\phi}$ at every date, both shares stay in $[\phi_0, \bar{\phi}]$ and converge together to the edge, $\phi_\infty = \bar{\phi}$.

A race between invention and adoption dictates which case obtains. Adoption (6) closes the gap $\phi^* - \phi$ at rate ξ , while entry widens it as long as $\iota_A > 0$, at a rate scaled by $\psi\theta_A$, so the outcome depends on the ratio $\psi\theta_A/\xi$, as shown in (42). With fast adoption relative to invention, $\psi\theta_A/\xi \rightarrow 0$, the installed share tracks the frontier, $\phi = \phi^*$, and entry stops at the edge case $\phi_\infty = \bar{\phi}$. With fast invention, the frontier already exceeds the floor $\bar{\phi}$ once the installed share reaches it, and an interior case $\phi_\infty = \bar{\phi}^*$ obtains, with $\bar{\phi}^*$ rising with the ratio $\psi\theta_A/\xi$. A critical value of this ratio then separates the two cases.

4 Simulation

We now turn to quantitative predictions of the model regarding long-run possibilities and the associated transitional paths. We start in Section 4.1 by describing model calibration, before turning to model simulation and scenarios in Section 4.2.

4.1 Calibration

To calibrate the model, we draw on both external anchors and judgements, as summarised by Table 1. Where credible measured quantities are available, we calibrate the parameters to

these benchmarks. In the absence of such empirical anchors, notably for some parameters governing AI’s potential, we resolve the uncertainty in AI’s favour so that any bottleneck cannot be attributed to the supply-side constraints.

Table 1: Calibration

	Parameter	Value	Rationale
ρ	discount rate	0.02	standard
σ	inverse IES	1	log utility
δ	depreciation rate	0.05	standard (Barro and Sala-i-Martin, 2004)
ε	demand elasticity	5	25% cost-weighted markup (Edmond et al., 2023)
α_A	AI share of product rents	0.40	pre-AI research bill of 4% of GDP
ϕ_0	initial automated share	0.40	labour factor share (BEA)
g_0^*	baseline trend growth	0.02	long-run US trend (Fernald and Jones, 2014)
$\iota_{H,0}$	baseline arrival rate	0.06	creative destruction (Garcia-Macia et al., 2019)
d_A	obsolescence, AI products	0.25	R&D obsolescence (Li and Hall, 2020)
d_H	obsolescence, conventional	0.04	consumer products (Broda and Weinstein, 2010)
$\iota_{A,0}$	AI arrival at activation	0.03	imputed from AI productivity estimates
$\phi_{\max}$	automatable frontier	0.95	Slack AI supply
ψ	frontier advance speed	5	Slack AI supply
ξ	adoption/research-access speed	0.35	Slack AI supply

Note: The table reports central parameter values used in simulation and their rationales in the last column. Rates are in annual continuous-time units. Derived quantities include the step size $\theta = 1/3$, the baseline consumption rate $\mu_0 = 0.690$, and the thresholds $\hat{\mu}_A = 0.621$ and $\hat{\mu}_H = 0.345$ of (43), all discussed in the text.

The first group of parameters is pinned to standard values in the literature. The discount rate of 2% is standard, the depreciation rate of 5% a year is from the growth literature (Barro and Sala-i-Martin, 2004).¹³ The log utility assumption implies $\sigma = 1$. The demand elasticity $\varepsilon = 5$ under CES implies a 25% markup, the recent level of the cost-weighted average markup of Edmond et al. (2023). The initial automated share, $\phi_0 = 0.40$, implies the labour task share at $1 - \phi_0 = 0.60$, consistent with the employee compensation averaging 62% of national income over 2015–2025 (BEA NIPA Table 1.12). Baseline trend growth $g_0^* = 0.02$ matches the long-run US trend (Fernald and Jones, 2014). The baseline arrival rate annualises the creative-destruction rates of Garcia-Macia et al. (2019) to their mid-point 0.06. These calibrations imply a pre-AI balanced-path return of $r = \rho + \sigma g_0^* = 4\%$, which is the model’s natural rate of interest. The gap between this and measured natural rates such as Holston et al. (2023), which are closer to 1%, is commonly interpreted as reflecting a safe-asset premium (Del Negro et al., 2017).

Obsolescence and AI arrival rates are less standard and are anchored to their closest

¹³Machines are generic, operating all newly automated tasks. To the extent that capital powered by AI is dominated by compute and software, it may depreciate faster. In this case, it can be shown that the consumer market thins faster, the bottleneck region grows wider and output growth falls below the pre-AI trend sooner. The baseline results are thus conservative.

proxies. The AI obsolescence rate draws on Li and Hall (2020), who estimate the decay of R&D’s private value across industries, ranging from 11% a year in pharmaceuticals to 49% in computer system design. We take $d_A = 0.25$ as the baseline value, conservative against the 30% mid-point of their range. For conventional obsolescence rate, the value-weighted product destruction rate of 3% a year from Broda and Weinstein (2010) provides one anchor. We set d_H slightly above it at 0.04, since a product’s rent erodes through imitation and competition before the product expires. The AI arrival rate has no close precedent and is imputed from estimated AI productivity impacts. We set $\iota_{A,0}$ so that the trend boost under the Bounded boost scenario (see below) is about half a percentage point a year, close to the mid-point estimate of Filippucci et al. (2024) (0.43 percentage points) and within the span of other studies, such as 0.06 in Acemoglu (2024) and 1.5 in Briggs and Kodnani (2023).

The diffusion parameters, also without standard anchoring, are set to be generous to AI potential. The automation ceiling $\phi_{\max} = 0.95$ adopts the view that nearly every task is eventually automatable. The ceiling is also the parameter that exposes the economy to the bottleneck: the further automation can run, the further the consumption rate can fall (Section 4.2). The speeds of frontier advance and adoption are deliberately set to be fast, with adoption closing the gap to the frontier with a mean lag just under three years ($\xi = 0.35$), much quicker than historical diffusion (Comin and Hobijn, 2010), and $\psi = 5$ implying that the frontier’s remaining gap to the ceiling halves for every 14% of cumulative capability growth. The same speed phases in research access after activation by (18). The two speed parameters affect only the transition path rather than the balanced-path objects themselves, and the trajectory depends on them only through the composite $\psi\theta_A/\xi$ (Section 3.3). We adopt these optimistic assumptions about AI potential and adoption to ensure that any bottleneck does not arise from the supply side.

The remaining model parameters are derived from these primitives. The innovation step size is $\theta = g_0^*/\iota_{H,0} = 1/3$ on both research dimensions, $\theta_A = \theta_H = \theta$. The AI share of product rents, $\alpha_A = 0.4$, splits each product’s rent between the two research dimensions. At this value the pre-AI research bill is 4% of GDP, above business spending on research and development at 2.6% of GDP (National Center for Science and Engineering Statistics, 2026), and well below total investment in intangibles, at over 15% (Corrado et al., 2009, 2022). The baseline consumption rate μ_0 is the resource relation (31) evaluated at ϕ_0 on the pre-AI balanced path, which gives 0.690, against a measured share of personal consumption in GDP of 0.68 in 2025 (Bureau of Economic Analysis, 2026). The anchored thresholds follow from free entry on the pre-AI balanced path. At the observed pair $(\mu_0, \iota_{x,0})$, the condition $V_x = \alpha_x \eta_x Y$ and the valuation (28) jointly pin the composite $\varepsilon \eta_x = \mu_0/(\rho + d_x + \iota_{x,0})$. Substituting the

composite into the threshold (30) cancels the unobservables and yields:

$$\hat{\mu}_x = \mu_0 \cdot \frac{\rho + d_x}{\rho + d_x + \iota_{x,0}}. \quad (43)$$

Evaluated at $(d_A, \iota_{A,0})$ and $(d_H, \iota_{H,0})$, the thresholds are $\hat{\mu}_A = 0.621$ and $\hat{\mu}_H = 0.345$.

The anchored thresholds imply arrival elasticities with respect to the consumption rate of 2.0 on the conventional margin and 10 on the nascent AI margin at activation. Acemoglu and Linn (2004) estimate an entry elasticity of 4 to 6 with respect to the potential size of a drug's market. Our conventional elasticity of 2.0 is below this range, as befits an economy-wide average, while the AI elasticity of 10 is justified by the young age of the sector with rarer arrivals and shorter-lived rents. We later consider the robustness of findings with respect to these values.

4.2 Scenarios

We compute the equilibrium path one period at a time. Under the quasi-static specification of Section 3.3, entrants price entry off the current consumption rate, so at each date the trend block solves (41) at the current ϕ_t and s_t . The capital stock evolves by (3), with investment determined by goods-market clearing (17), and the frontier and the automated share by (5)–(6). Trend growth g_t^* is the balanced-path rate at the current automated share, while realised growth g_t is the growth of output along the simulated path, along which the capital stock is being built up and tasks being reassigned. The gap $g_t - g_t^*$ is negative during the initial phase of capital buildout, turns positive as the capital stock catches up, and converges to zero as the automated share settles (Figure 4). All simulation results are shown in years since activation t_0 .

We adopt a temporary equilibrium concept with static expectations: at each date agents optimise and markets clear under the belief that current conditions persist. An alternative concept is one of perfect foresight, where agents fully anticipate the dynamic adjustments of the AI transition. The temporary equilibrium concept is arguably more realistic, given the difficulties of forecasting the implications of a once-in-a-lifetime transformation decades ahead. Abstracting from anticipation also keeps the mechanism transparent, with entry responding to observed spending alone. We consider the perfect-foresight case at the end of this section, showing that all key qualitative results are preserved.


We distinguish five scenarios in our simulation, assuming parameter values as in Table 1 but varying the more uncertain AI-related parameters.

1. **Baseline** assumes no AI activation, or equivalently an activation without arrivals due to supply-side bottlenecks in the sense of Aghion et al. (2019), Trammell and Korinek (2026) and Jones and Tonetti (2026) keeping AI inert.
2. **Bounded boost** assumes activation with at most limited expansion of the frontier and new automation. The drawn path takes the limiting case of $\phi_{\max} = \phi_0$, so AI deepens capability on the tasks machines already perform and automates nothing new.
3. **Transformative** assumes near-immortal rents and a fast arrival, with $(\iota_{A,0}, d_A)$ that tracks the supply-side benchmark as closely as the model allows (Appendix B.2).
4. **Automation without innovation** assumes an intermediate combination of AI arrival and rent durability: strong enough that automation runs to the frontier, yet too weak to sustain the innovation engine. The drawn path takes the case in which AI innovations arrive and their rents decay at the same rates as conventional innovations.
5. **Demand bottleneck** assumes the central calibration.

Figure 3 shows the simulation paths of the five scenarios. *Bounded boost* lifts trend growth and the natural rate permanently but modestly, by about half a percentage point, through sustained AI innovation. The labour market impact is benign, with the task share unaffected at 60% and real wages rising throughout. *Transformative* brings a continued increase in trend growth, reaching 6% by the end of the simulation, and a corresponding rise in the natural rate. The AI arrival rate rises to a peak several years after activation and then declines, but growth keeps rising as each capability improvement raises output on a growing share of tasks. The labour task share collapses to under 10%. Wages initially fall below their no-AI path, as automation displaces tasks faster than output grows, but eventually recover as the remaining human tasks command a rising premium. *Automation without innovation* automates most tasks but its AI innovation fades, so trend growth rises only temporarily, declines after a few years and ends below the Baseline. Under *Demand bottleneck*, trend growth and the natural rate rise briefly and modestly but fall through the Baseline within a decade as AI arrivals quickly evaporate. The labour task share stalls at 48% as the consumer market thins. The natural rate settles below its pre-AI level.

Real wages fall relative to the Baseline in every scenario with automation, and whether they recover depends on a race between output and the task share. Automation takes tasks from workers immediately, while the extra output arrives only as the machines for those tasks are built. Since the wage per worker is $(1 - \phi_t)Q_t$ by (14), its gap to the Baseline path widens first and starts to close only once output outpaces the Baseline by more than the task

Figure 3: Simulation results under five scenarios



Note: Simulated paths under the five scenarios in years since activation. Panel (a): trend growth g_t^* with the natural rate $r_t^* = \rho + \sigma g_t^*$ on the right axis. Panel (b): the AI arrival rate $s_t \iota_{A,t}$. Panel (c): the labour task share $1 - \phi_t$. Panel (d): the real wage per worker $(1 - \phi_t)Q_t$ relative to its Baseline path $(1 - \phi_0)Q_t^0$.


share shrinks, $g_t - g_0^* \Rightarrow \dot{\phi}_t / (1 - \phi_t)$. *Bounded boost* displaces no one, so its wage rises with output from the start. Under *Demand bottleneck* the task share loss is permanent and trend growth ends below the Baseline, so the wage settles about 15% below the Baseline path. The two scenarios that automate to the frontier cut the task share by half or more within two decades, and leave real wage roughly 50% below the Baseline path. Under *Transformative* the real wage eventually recovers as faster output growth makes up for the falling labour task share: the few remaining human tasks are the scarce complement to machine-produced output. Under *Automation without innovation* the recovery is slower, and the wage is still 40% below the Baseline path five decades after activation. These paths assume a fixed labour supply, no new tasks for labour and wages at the marginal product, which is what makes the displacement permanent.

Output growth has four sources: conventional innovation, AI innovation, the reassignment of tasks to machines and capital deepening (Figure 4). Trend growth is the sum of the first two, representing the innovation sustained by the consumer market at the current automated share. Task reallocation subtracts from growth initially when automation is fast relative to the capital buildout. Capital deepening then adds to growth once the capital stock catches up, and keeps adding while the stock grows faster than trend. Under *Bounded boost* no task is reassigned and output growth tracks trend growth. Under *Demand bottleneck* the AI contribution is largest in the first years after activation and vanishes when entry stops, reallocation subtracts up to a fifth of a percentage point around the end of the first decade, and capital deepening adds almost a percentage point in the second decade. Output growth therefore peaks just under 2.4% about 15 years after activation, when trend growth is already below the Baseline, and falls below the Baseline after about 35 years. In *Automation without innovation*, capital deepening adds over two percentage points for multiple decades, so output growth stays above the Baseline throughout the window. In *Transformative*, AI innovation replaces conventional innovation entirely within three decades, as the thinning consumer market shuts entry on the conventional dimension while near-immortal AI rents sustain AI entry.

A key common mechanism separates these outcomes. Every newly automated task must be equipped with capital, which must be grown and replaced thereafter. Automation thus permanently raises the investment share of output and lowers the consumption rate. But innovation is paid out of consumer spending through the product rent $\mu Y/\varepsilon$, while machines are competitively priced and repay no research. A thinning consumption rate therefore lowers the prize for both AI and conventional innovation. The scenarios differ only in how much of the prize survives. With no new automation the thinning is limited, which is the *Bounded boost* case. In *Transformative*, near-immortal rents keep the prize large even at a compressed consumption rate, arrivals keep coming, and growth takes off. In *Automation without innovation*, rents last only as long as the conventional technology's, so the prize shrinks with the consumption rate yet never drops below the entry cost. Developing new AI products always pays a little, but each further step of automation leaves fewer worthwhile arrivals and growth fades rather than stops. In *Demand bottleneck*, rents are short-lived and the prize falls to the entry cost early: μ falls from 0.690 through the anchored threshold $\hat{\mu}_A = 0.621$, a 10% thinning, and entry into new AI products stops almost two decades after activation. Conventional innovation continues but weakens with the same spending. ι_H falls from 0.060 to 0.048, the product of a 10% permanent thinning and a conventional arrival elasticity of two, taking trend growth from 2.0% to 1.6%.

Figure 5 panel (a) shows how far the consumer market can thin before AI entry stops being

Figure 4: Sources of output growth




Note: Decomposition of output growth ($\dot{Q}_t/Q_t$) into four components: conventional innovation θ_{ι_H} , AI innovation $\theta_{s_{\iota_A}} \phi_t / (1 - \phi_t)$, task reallocation $\phi_t [\ln(A_t K_t / \phi_t) - \ln(H_t / (1 - \phi_t))]$ and capital deepening $\phi_t (g_K - g^*)$. Panels correspond to the four AI scenarios. Hatched areas are negative contributions. The solid line is output growth and the dashed line trend growth, the sum of the two innovation components. Dotted rules mark zero and the pre-AI 2%.

worthwhile. The AI entry threshold differs across scenarios. It is lower the longer AI rents last, since a longer-lived rent can repay entry even from a thinner market. It is also lower the higher the arrival rate at activation. Baseline, Bounded boost and Demand bottleneck share the same threshold (shown in dash-dot blue line). Automation without innovation, where AI rents last only as long as conventional rents, has a far lower threshold. Transformative, where AI rents are near-immortal, has the lowest threshold. Under Bounded boost the consumer market hardly thins. Under Demand bottleneck the consumption rate reaches the threshold and halts entry in less than two decades. Under the other two scenarios the consumption rate falls much further, but the threshold lies further away, so entry continues throughout the window. Panel (b) shows the same distance as the share of the pre-AI entry rate the

consumer market still supports, which reaches zero only under Demand bottleneck.

Figure 5: The consumer market and AI entry




Note: The consumer market and AI entry under the five scenarios in years since activation. Panel (a): the consumption rate μ_t , with each scenario's AI entry threshold $\hat{\mu}_A$ as a dash-dot line. Panel (b): the AI arrival rate that the current consumption rate can support, as a share of its pre-AI level, $(\mu_t - \hat{\mu}_A)/(\mu_0 - \hat{\mu}_A)$ in per cent, clipped at zero.

The parameters dictate which of the outcomes materialises. Figure 6 shows the parametric regions corresponding to the scenarios, starting with panel (a) which depicts the existence of Transformative, Automation without innovation and Demand bottleneck scenarios as a function of AI arrival rate at activation $\iota_{A,0}$ and the AI rent obsolescence rate d_A . All other parameters are held fixed at their central values in Table 1. The top boundary traces the combinations $(\iota_{A,0}, d_A)$ at which the balanced-path consumption rate at full automation just equals the AI threshold: above this boundary automation stalls prematurely, resulting in demand bottleneck. The lower boundary traces the pairs at which the fully automated economy grows at the pre-AI 2.0%, below which productivity takes off and the economy attains transformative AI. The band between the two boundaries is the Automation without innovation region.

Imposing empirical discipline on the parameter ranges sheds light on the relative likelihood of these three scenarios. Direct AI-specific evidence on $(\iota_{A,0}, d_A)$ remains lacking, but the closest existing evidence discussed in Section 4.1 places the economy in the Demand bottleneck region. Obsolescence d_A ranges between 0.11 and 0.49 a year according to Li and Hall (2020). The arrival rate $\iota_{A,0}$ is inverted from estimates of AI's boost to productivity growth, which run from 0.06 to 1.5 percentage points a year (Acemoglu, 2024; Briggs and Kodnani, 2023),

Figure 6: Scenario regions



Note: Scenario regions in the space of AI rent obsolescence d_A , AI arrival rate at activation $\iota_{A,0}$, and automatable frontier $\phi_{\max}$, with all other parameters fixed at their central values. Panel (a) varies $\iota_{A,0}$ and d_A at the calibrated frontier $\phi_{\max} = 0.95$. Panel (b) varies $\phi_{\max}$ and d_A at the central arrival $\iota_{A,0} = 0.03$. The black line is the stall condition from (43), $\hat{\mu}_A(d_A) = \mu(\phi_{\max})$, above which the steady-state interval is non-empty. The grey line is where the balanced path at the automation ceiling grows at the pre-AI 2.0%. Between the black and grey lines, the automation without innovation region, automation completes while steady-state trend growth lies at or below 2%. The navy dots mark the central calibration. The purple square and the amber circle mark the Automation without innovation and Transformative scenarios of Figure 3, and the red circle denotes the Bounded boost scenario.

implying a range for $\iota_{A,0}$ of 0.003 to 0.084.¹⁴ The rectangle defined by the two intervals lies inside the bottleneck region, well above the top boundary, which reaches only $d_A = 0.074$ even at the rectangle's right edge. At that edge, AI rents would need to be about one and a half times as durable as the most generous empirical benchmark for the economy to escape the bottleneck region, and nearly three times as durable to enter the transformative region.

In panel (b) of Figure 6, we sweep the automation frontier $\phi_{\max}$ and obsolescence d_A , bringing the Bounded boost scenario into view. All other parameters are held at their central values in Table 1. Again, two boundaries divide the plane. Along the central obsolescence $d_A = 0.25$, frontiers up to 0.47 deliver the bounded boost to productivity growth. Frontiers between 0.47 and 0.51 deliver automation without innovation, and beyond 0.51 the demand bottleneck sets in. The bottleneck region widens as the automatable frontier expands, and above $\phi_{\max} \approx 0.61$ it covers the entire estimated obsolescence interval of 0.11–0.49, placing

¹⁴The inversion uses the Bounded boost scenario, in which the automated share stays at ϕ_0 . On the balanced path at that share, (41) makes the productivity boost an increasing function of the arrival rate alone. Inverting 0.06 and 1.5 then gives $\iota_{A,0} = 0.003$ and 0.084.

the central calibration deep inside the bottleneck region. Escaping the bottleneck at the central obsolescence requires automation to stop within about 10 percentage points of the current share.

Within the bottleneck region, the obsolescence rates d_A, d_H and the AI arrival $\iota_{A,0}$ also govern the bottleneck’s severity and how fast it sets in. We re-run the simulation varying one parameter at a time while holding the others constant at central values, and assess the steady-state shortfall in trend growth, which is -0.41 percentage points in the central calibration. The AI arrival rate is the most consequential: varying $\iota_{A,0}$ across 0.003 – 0.084 moves the shortfall from -0.05 to -0.95 . Sweeping AI obsolescence across its empirical range of 0.11 – 0.49 puts the shortfall between -0.75 and -0.26 , while adjusting conventional obsolescence across 0.02 – 0.06 moves it between -0.35 and -0.48 . Across all 32 combinations of the five parameters at the edges of their ranges, the steady-state shortfall spans -0.02 to -1.83 and the steady-state labour task share 13% to 59% (Appendix B.3). The two diffusion speeds ψ and ξ , the latter moving both the adoption lag and the phase-in of research access, each move the shortfall on their own by less than a tenth of a percentage point. In every combination examined, trend growth ends below the pre-AI baseline.

Incorporating perfect foresight leaves the entry threshold and the set of steady states unchanged and moves the steady state very slightly by changing how AI entry ends. Solving the full perfect-foresight equilibrium, with capitalists on their Euler equation and free entry holding along the actual path (Appendix B.1), lowers the boom peak from 2.32 to 2.25% and brings the crossing of the pre-AI rate forward from year 9 to year 5, because capitalists who foresee the thinning consumer market cut consumption and invest more. AI entry then fades instead of stopping, and the economy settles at the lower edge of the steady-state interval, with the automated share at 0.514 rather than 0.518 and steady-state trend growth at 1.60 rather than 1.59% .

5 The AI debate

In this section we place our findings in the broader AI debate. We start by comparing the model’s growth projections with those of recent studies and examine the sources of the differences. Second, we take a closer look at the labour market and distributional implications of the model. Third, we discuss the implications for monetary policy. Finally, we evaluate the role of other policies in securing the gains from AI.

5.1 Comparison of growth scenarios

The first debate is how much transformative AI could add to growth. Korinek and Suh (2024), Korinek et al. (2026) and Jones and Tonetti (2026) provide quantitative scenarios for economic growth following transformative AI.¹⁵ Table 2 summarises the results of these studies together with our projections. To aid comparison, we align the studies’ predictions by taking the start of each simulation as the activation date.¹⁶ Korinek et al. (2026) initialise the simulation at mid-2026 and project the state of the economy three and a half years later, ending at the start of 2030, which we take as equivalent to year 4 of our simulation. Jones and Tonetti (2026) simulate from 2020 and report 2030, 2040 and 2050, which we equate to our years 10, 20 and 30. Korinek and Suh (2024) index time from the date at which automation starts to advance and place artificial general intelligence, the automation of every task, at year 20 in their baseline scenario and at year 5 in their aggressive one.

Our simulations generally point to lower growth outcomes, even under the more optimistic Transformative scenario. At year 4, all our scenarios produce output and growth predictions in the range of others’ most modest scenarios. The gap widens further with the horizon. By year 20, Jones and Tonetti (2026)’s Moore’s law case grows at 13% with output three times its trend, while Korinek and Suh (2024)’s baseline scenario grows at 18%. In our case, growth reaches 5.5% and output exceeds the no-AI path by 39% under the Transformative path, and only 2.3% and 6% on the Demand bottleneck path.

Three factors account for the substantial differences between our estimates and other studies’. The first is the assumed nature of transformative AI. While all studies associate it with more automation, the trigger differs. Korinek and Suh (2024) assume that automation advances at an exogenous rate in the manner of Moore’s law, Jones and Tonetti (2026) set the automation rate at the computer sector’s historical pace in their Moore’s law calibration, Korinek et al. (2026) build scenarios from five exogenous paths including the automation share, and our model endogenises the research decision that advances the frontier for as long as entry is profitable. The second factor is calibration. Other studies assume a relatively rapid automation under optimistic scenarios, covering a quarter of tasks within four years and all tasks within two decades. In our case, automated shares increase after two decades by only 12 percentage points under Demand bottleneck and by 40 percentage points under

¹⁵Other studies put the gain from AI much lower: Acemoglu (2024) at no more than 0.66% of total factor productivity over ten years, and Filippucci et al. (2024) at 0.25 to 0.6 percentage points of annual TFP growth. Both hold the set of exposed tasks fixed at today’s exposure estimates, about 5% of tasks in the first case, and rule out changes in the process of invention within their ten-year horizon. As they do not consider transformative AI in our sense, we leave them out of the table.

¹⁶The alignment is not exact. Other studies’ starting dates are calendar dates at which AI already diffuses, while ours is the date from which research can widen the automatable frontier.

Table 2: Growth scenario comparisons

	Output growth, %	Output vs no-AI, %	Wage vs no-AI, %	Labour share, %
Short term (4 years)				
Korinek et al. (2026), modest	2.4	1.6	0.7	59.4
Korinek et al. (2026), substantial	5.4	8.3	2.1	56.1
Korinek et al. (2026), extreme	15.4	32.4	9.7	45.2
Korinek and Suh (2024), aggressive AGI ^a	18.0			→0
This paper, Bounded boost	2.3	0.7	0.7	60.0
This paper, Demand bottleneck	2.3	0.7	−2.1	58.3
This paper, Automation without innovation	2.6	1.4	−4.3	56.6
This paper, Transformative	2.8	1.9	−6.0	55.3
Long term (20 years)				
Jones and Tonetti (2026), Moore’s law	13.0	220 ^b		62 ^c
Korinek and Suh (2024), baseline AGI	18.0			→0
This paper, Bounded boost	2.5	7.4	7.4	60.0
This paper, Demand bottleneck	2.3	6.2	−14.5	48.3
This paper, Automation without innovation	3.9	21.7	−43.9	27.6
This paper, Transformative	5.5	38.9	−54.2	19.8

Note: Other studies’ simulation start is matched with our activation date. Horizons are chosen to align the studies as much as possible. Year 4 (short term) is where Korinek et al. (2026) stop, with projections taken from their Table 3. Year 20 (long term) is 2040 in Jones and Tonetti (2026) who simulate from 2020. Their projections are taken from Figure 6. Korinek and Suh (2024) do not report a no-AI path. Other studies’ labour share is a share of income, while ours is the task share $1 - \phi_t$. Our entries come from the simulation of Section 4.2, against a Baseline path with growth of 2.0% and a task share of 60% throughout. ^a Reaches full automation in year 5, reported in the four-year group. ^b Relative to their own initial trend, which grows at 4.7% a year in this calibration. ^c Extracted from chart.

Transformative path. When the assumed paces of automation are more similar, as in the modest scenario of Korinek et al. (2026), growth simulations match more closely. The third, and most consequential, factor is the framework for ideas and capital. In other studies, research effort is a fixed share of output or an exogenous endowment, and capital arrives along a schedule, as a constant investment share or, in Korinek and Suh (2024), through optimal saving with automation advancing on an exogenous schedule. The pace of innovation does not respond to how much of output is consumed. Our framework requires that research be paid for out of consumer spending, a more stringent condition for sustaining growth. Once automation thins the consumer market, the growth engine can stall.

5.2 Labour market and income redistribution

The second debate is what consequences AI has for the labour market and for inequality. As Table 2 shows, most studies predict a falling labour share as automation proceeds, with

differences in the pace and the extent of the fall and in the implications for the wage. Korinek et al. (2026) predict an increase in the real wage by 1 to 10% in four years even as the labour share falls from 60 to 45% in the extreme scenario, because each of their AI paths raises the productivity of the remaining labour tasks and capital accumulates freely along a schedule. Our wage is the task share times output, with output rising slowly because the machines are built out of saving. Hence the wage falls behind the no-AI path, by 2% at year 4 and 15% at year 20 on the Demand bottleneck path, and by more than half at year 20 on the Transformative path. For the labour share, Jones and Tonetti (2026) predict a stable share near 60% because tasks are strong complements, while Korinek and Suh (2024) project a share of zero under full automation. Our estimates span 14 to 48% across the scenarios that automate within three decades, depending on whether entry stalls or survives.¹⁷


An important assumption underlying recent studies and this paper is that transformative AI goes hand-in-hand with more automation.¹⁸ What if labour retains its critical function and works alongside more capable machines in a complementary manner? In our Bounded boost scenario, machines improve at the tasks they already perform and replace no further tasks, but the capability gains still accrue to machines alone. Workers' productivity remains unchanged and their wage rises only because tasks are complements. We can relax this assumption and let each advance in machine capability raise workers' productivity at their tasks. So far we have ruled out this potential spillover from AI research into labour productivity by assigning labour-augmenting progress to conventional innovation alone.

Figure 7 shows what this productivity spillover does. When applied to the Bounded boost case, where there is no task displacement, trend growth settles at 3.3% rather than 2.5%, the wage rises to two-thirds above the no-AI path rather than a quarter, and the task share stays at 60%. When applied to the Demand bottleneck path, the same spillover amplifies the short-term boom, delays the stall of AI entry by two years and leaves the steady state all but unchanged since it does not alter the capital deepening that thins the consumer market. Complementarity therefore adds to steady-state growth only when it comes without labour displacement. For the labour share to hold steady while growth exceeds Bounded

¹⁷Two technical issues bear on this comparison. Other studies' labour share is a share of income, while ours is a share of tasks. The income share equals the task share times one minus the profit share of GDP, since wages are paid out of output at cost. Other studies' predicted unemployment, 3 to 18% of cognitive workers by 2030, stems from a search friction, which in our full-employment framework would show up in the wage instead.

¹⁸In our model this is the definition: at activation, AI research can keep improving machine capability and widening the set of tasks machines perform. Korinek and Suh (2024) define artificial general intelligence as the ability to perform every task humans perform, Jones and Tonetti (2026) describe transformative AI as automating all cognitive tasks, and the extreme scenario of Korinek et al. (2026) has AI performing almost half of today's cognitive work by 2030.

Figure 7: Complementarity between workers and machines



Note: Simulations with and without a spillover from AI research to workers' productivity, under the Bounded boost and Demand bottleneck scenarios. Panel (a): trend growth g^* in per cent. Panel (b): the real wage per worker $(1 - \phi_t)Q_t$ relative to its Baseline path, with the dotted rule at 1. Panel (c): the labour task share $1 - \phi_t$ in per cent. Solid red and navy are the Bounded boost and Demand bottleneck paths of Figure 3, and black dashed is the Baseline. Dashed red and navy add the spillover, under which each AI advance raises workers' productivity at the tasks they keep by as much as it raises machine capability. Dotted navy adds instead a one-off gain of a fifth in labour productivity at activation. All other parameters are at their calibrated values.


boost, what is needed is not a one-off gain in labour productivity but AI research that keeps improving machine and worker productivity together, without widening the set of tasks machines perform. Another route, which the model does not incorporate, is the reinstatement of new tasks in Acemoglu and Restrepo (2019): automation proceeds but new tasks for labour are created at the same pace, so the net task share holds while each automated task keeps getting cheaper. Adding it would mean letting the set of tasks grow with automation, and the consumption rate would then respond to the net automated share.

5.3 Monetary policy implications

The third debate concerns the implications of transformative AI for monetary policy, in particular for the natural rate of interest, or r -star. The canonical hypothesis is that, to the extent that AI lifts productivity growth, r -star should increase. Our simulations in Section 4.2, however, show the wide range of r -star paths that could materialise across the scenarios. This scenario uncertainty comes on top of the well-known parameter and model uncertainties associated with r -star estimates.

Figure 8 shows the probability distribution of the natural rate 20 years after activation obtained by varying the primitive parameters governing the transition across their empirical

ranges. We draw 2,000 combinations of the five key parameters uniformly from the ranges of Appendix B.3, together with the rent-base share λ from zero to a conservative upper bound. For each combination we solve the transition and record the trend natural rate, generating the resulting distribution as the thick curve. The two thin curves show the distribution conditional on a fixed rent base. At $\lambda = 0$ innovators earn rents on household purchases alone, the main calibration. At $\lambda = 1$ innovators earn rents on all spending and AI entry never stops. The distribution is illustrative as it treats every combination in the ranges as equally likely. We normalise the model's natural rate with a level shift, $\hat{r}^* = r^* + z$, such that the pre-AI $\hat{r}^*$ equals the latest Laubach-Williams estimate of r^* for the United States, about 1.7%. After the activation date we allow z to be stochastic, with a standard deviation of 1.3 percentage points, the sample-average standard error of the Laubach-Williams estimate, to capture the typical uncertainty of the estimate.



Note: Probability densities of the natural rate $\hat{r}^*$ at year 20 after activation. λ is the fraction of firms' spending on machines and research inputs on which innovators earn the markup. At $\lambda = 0$, the paper's calibration, innovators earn rents on household purchases alone and AI entry eventually stops. At $\lambda = 1$, rents come from all spending and entry never stops. Thick curve: 2,000 combinations of $\{d_A, d_H, \iota_{A,0}, \psi, \xi, \lambda\}$ drawn uniformly from the empirical ranges of Appendix B.3 and $\lambda \in [0, 0.66]$, the upper limit being the share of firm purchases from the AI-producing industries that adds to the capital stock (2023 BEA input-output and fixed-asset tables). Thin curves: the same draws with λ fixed at 0 and at 1. For each draw, $\hat{r}^*$ is the trend natural rate shifted by the constant that maps the model's pre-AI rate of 4% onto the Laubach-Williams estimate for the United States of 1.65% (Laubach and Williams, 2003; Holston et al., 2023), and each curve averages normal densities centred on the draws with a standard deviation of 1.28 percentage points, the average standard error of the Laubach-Williams estimate (as of August 2026). The percentiles quoted in the text are those of these curves. The horizontal axis ends at 12%, beyond which lie 29% of the mass at $\lambda = 1$. The vertical dashed line marks the pre-AI level.

The range of the natural rate is wide but its median is close to the pre-AI level. At $\lambda = 0$, the main calibration, the median natural rate is 1.2%, half a point below the pre-AI level. With the rent base drawn from zero to two thirds, a conservative upper bound, the median is 1.6% and the 90% band spans -0.6% to 4.4% . Only when the markup extends to all spending does the distribution shift and acquire a long right tail, with a median of 6.7% and a 90% band reaching 24%, because AI entry then never stops. The evidence does not point to the right tail. Uncertainty is very high, however. The estimation error alone spans four points, so it remains difficult to draw a definitive conclusion about the direction of travel for r -star, let alone a precise quantitative estimate.

5.4 Broader policy implications

The final debate is the role of policy in securing the gains from AI and minimising its unintended effects.

From our model’s perspective, mitigating the demand bottleneck is imperative to maximise the growth impact of AI. One possible lever is the cost of research. A permanent cut by 10% in the cost of conventional research, for example through subsidised use of AI tools by ordinary researchers, returns trend growth to its pre-AI 2% after the stall, because cheaper conventional research entry replaces the AI one that is no longer worthwhile. Averting the bottleneck itself requires keeping AI entry worthwhile. Sufficiently cheap AI research could do so and nudge the economy onto the Transformative path but leave labour performing a smaller share of tasks (Appendices C.1 and C.2). Another lever is the innovator’s prize, the mirror image of the first. A tax on rents at rate τ acts like a rise in the cost of research by the factor $1/(1 - \tau)$. Policies that lengthen the life of innovators’ rents, such as stronger protection of intellectual property, work similarly, but at the calibrated arrival rate AI rents would have to last implausibly long for entry to survive to full automation, and even perpetual rents would not return trend growth to 2%. On the other hand, redistributive policy that shifts consumption between households would be ineffective as a growth lever, since the entry decision is driven by the share of output consumed rather than the identity of consumers.

Maximising growth alone may not be the policymaker’s goal, however, as the consequences for the labour market may well be unacceptable. Instead, the policy objective may be to secure the productivity gains without incurring the labour-market cost. This means steering AI toward complementing workers rather than replacing them. The Bounded boost scenario is one example, where AI improves machine productivity without taking tasks away from workers. Another possibility is the complementarity extension above where AI research improves workers’ and machines’ productivities concurrently, keeping the task share stable.

The policies that deliver this must re-orient AI research toward complementing workers’ production rather than replacing them, in the spirit of Acemoglu and Johnson (2023) and Brynjolfsson (2022).

Our model abstracts from a number of important considerations. One is the security implications of advanced AI systems. Transformative AI has the potential to cause large-scale disruptions, possibly with impacts on long-term growth and beyond. Even before any damage occurs, a growing awareness of such risk might raise the discount rate, pushing up entry thresholds and constraining growth. We also abstract from supply-side bottlenecks, whether from essential tasks that resist automation in the sense of Aghion et al. (2019) and Jones and Tonetti (2026), or from limits in compute and energy, which could bind before the demand bottleneck does and protect the consumer market from thinning.

6 Conclusion

The debate over transformative AI and its growth impact has to date focused on supply-side forces: whether AI can propel growth via productivity and automation or will be held back by critical tasks that remain with humans. In this paper, we study a bottleneck that originates from the demand side, arising when the consumer market becomes too thin to reward further innovation and automation. The framework nests this new bottleneck possibility alongside the transformative one and the scenarios in between. When the model is calibrated to the available evidence, the demand bottleneck scenario emerges as highly plausible. The result argues for more tempered expectations of what AI can deliver for growth, even when the technology itself is transformative. The large-scale automation enabled by AI this time around could erode the consumption base and weaken the profit incentive that has spurred technological progress up to now. Transformative AI may thus carry the seed of its own slowdown. These conclusions rest on the assumptions that rents are earned primarily in the consumer market and that automation is rapid enough to outpace any new tasks created for labour. While plausible, both remain uncertain and a wide range of outcomes for growth, labour markets and the natural rate of interest is possible. Our framework nonetheless highlights the crucial role of market incentives in realising the technological promise of AI, incentives that may themselves be disrupted by the very progress they make possible.

Appendix A: Microfoundation

A.1 Aggregate production function

Real output in (1) is a geometric mean over products of geometric means over tasks. We derive it from the product level, with the rental $r_t + \delta$ and the wage w_t per worker common to all products, and show that the heterogeneity of products goes into prices and quantities rather than inputs.

Product j is made from a unit continuum of tasks. Machines perform the automated tasks, measured by ϕ_t , each with capital $k_{j,i}$ producing $A_{j,t}k_{j,i}$. Workers perform the rest, each task with $n_{j,i}$ workers producing $H_{j,t}n_{j,i}$. Tasks combine with unit elasticity, so product output is the geometric mean of task outputs. Because the log is concave, the cost-minimising allocation spreads each factor evenly across its tasks, $k_{j,i} = K_{j,t}/\phi_t$ and $n_{j,i} = L_{j,t}/(1 - \phi_t)$, and

$$q_{j,t} = \left(\frac{A_{j,t}K_{j,t}}{\phi_t} \right)^{\phi_t} \left(\frac{H_{j,t}L_{j,t}}{1 - \phi_t} \right)^{1 - \phi_t}, \quad (\text{A.1})$$

which is (1) with the product's indices and inputs, the form of Zeira (1998) and, in the unit-elasticity limit, of Aghion et al. (2019). Their Cobb–Douglas special case and Acemoglu and Restrepo (2018) absorb the denominators into a constant, which is exact on a balanced path when the task share is fixed. We retain the denominators because the task share ϕ_t moves along the transition in our model. Minimising $(r_t + \delta)K_{j,t} + w_tL_{j,t}$ for a given output makes each factor bill a fixed share of the cost outlay and the unit cost $c_{j,t}$ a geometric mean of the factor prices per efficiency unit:

$$(r_t + \delta)K_{j,t} = \phi_t c_{j,t}q_{j,t}, \quad w_tL_{j,t} = (1 - \phi_t) c_{j,t}q_{j,t}, \quad c_{j,t} = \left(\frac{r_t + \delta}{A_{j,t}} \right)^{\phi_t} \left(\frac{w_t}{H_{j,t}} \right)^{1 - \phi_t}. \quad (\text{A.2})$$

Every buyer spends the same amount on every product, because consumers, investing firms and researchers combine products with the same unit-elastic composite. Consumers pay the markup $\varepsilon/(\varepsilon - 1)$ over unit cost, and the other buyers pay cost. The cost outlay on a product is therefore the same for every product,

$$c_{j,t}q_{j,t} = \frac{\varepsilon - 1}{\varepsilon} C_t + I_t + R_t \equiv X_t, \quad (\text{A.3})$$

where the first term is consumer spending net of the markup. By (A.2), every product then employs the same capital and labour,

$$K_{j,t} = \frac{\phi_t X_t}{r_t + \delta} = K_t, \quad L_{j,t} = \frac{(1 - \phi_t) X_t}{w_t} = L_t, \quad (\text{A.4})$$

the second equalities because products have unit measure and the labour supply is one. The indices appear in quantity and price alone: $q_{j,t} = X_t/c_{j,t}$, and a product with a higher index sells more units at a lower price from the same inputs.

Real output is the geometric mean of product outputs, $\ln Q_t = \int_0^1 \ln q_{j,t} dj$, so with common inputs the indices aggregate into geometric means as well. Substituting (A.1) with the common inputs (A.4),

$$\ln Q_t = \phi_t \ln \frac{K_t}{\phi_t} + \phi_t \int_0^1 \ln A_{j,t} dj + (1 - \phi_t) \ln \frac{1}{1 - \phi_t} + (1 - \phi_t) \int_0^1 \ln H_{j,t} dj, \quad (\text{A.5})$$

which is (1) with $\ln A_t = \int_0^1 \ln A_{j,t} dj$ and $\ln H_t = \int_0^1 \ln H_{j,t} dj$. The aggregation is exact at every date, and no assumption on the distribution of the indices enters. The unit cost of the composite aggregates the same way: $\ln P_t = \int_0^1 \ln c_{j,t} dj$ gives $P_t = ((r_t + \delta)/A_t)^{\phi_t} (w_t/H_t)^{1-\phi_t}$ with the same averaged indices. Real output at cost is the numeraire, $P_t = 1$, so $X_t = Q_t$ and (A.4) gives the factor payments $(r_t + \delta)K_t = \phi_t Q_t$ and $w_t = (1 - \phi_t)Q_t$, the pricing of (14) with the wage written per worker rather than per efficiency unit.

Every incumbent earns the same profit flow. Consumer spending on each product is C_t and serving it costs $C_t(\varepsilon - 1)/\varepsilon$, so the incumbent earns

$$\pi_{j,t} = \frac{C_t}{\varepsilon} = \Pi_t, \quad (\text{A.6})$$

for every j , which is (13). An innovation raises the innovator's output and lowers its price, and its private return is the incumbency it takes over.

A.2 Firms' automation decision

The main text assumes that firms automate tasks at a fixed speed once machines can perform them. This appendix lets firms decide whether and how fast to automate, and shows that the decision delivers the rule's assumption, namely that firms keep automating after AI innovation stops and never switch back.

Let competitive task producers choose which tasks to automate and when. Converting a task from labour to machines yields the producer the present value of the task's factor-cost saving s_t , and installing at aggregate rate $\dot{\phi}_t$ costs $c(\dot{\phi}_t)$, an increasing and convex function, so the producer converts until the marginal installation cost equals that value,

$$c'(\dot{\phi}_t) = \int_t^\infty e^{-\int_t^\tau r_u du} s_\tau d\tau. \quad (\text{A.7})$$

Two properties of the saving are key. First, s_t is positive for every task between the installed share ϕ_t and the frontier ϕ_t^* , because a task inside the frontier runs cheaper on machines than on labour at prevailing factor prices. Second, the saving does not depend on the consumer market, as installed machines sell their services competitively at the rental $r_t + \delta$ whatever happens to innovation and the consumption rate μ_t . Producers therefore continue to automate after AI entry stops, until the installed share reaches the frozen frontier $\bar{\phi}^*$, as under the rule, at a speed set by the installation cost in place of ξ . Switching tasks back to labour never pays, since on the stalled path effective capital and effective labour grow at the same rate g , so relative factor prices are constant and the installed assignment remains the cheapest.

A.3 Innovators' choice of rent base

The main text assumes that innovators earn their rents from a markup on a consumer product. The alternative is to let innovators earn rents from investment instead, by charging licence fees on firms that install the machine. This appendix considers the innovators' choice and shows that it depends on the longevity of rents on machines, so the assumption can be confronted with evidence.

Let each AI entrant pick one of two bases at entry. A consumer incumbency pays the flow $\mu Y/\varepsilon$ and ends at hazard $d_A + \iota_A$. Selling the innovation to the firms installing automation pays the flow $(I/Y)Y/\varepsilon$, and ends at hazard $d_B + \iota_A$. The entry cost and the markup parameter are the same on both bases. The two bases differ in only the size of the payoff streams and the rent longevity, with d_B being the new parameter.

An incumbency is worth its payoff stream divided by ρ plus its hazard, and entry pays when that value covers the entry cost. On the consumer base, Section 4.1 anchors the AI entry cost η_A by imposing free entry on the pre-AI balanced path, which at the calibrated μ_0 and $\iota_{A,0}$ implies $\varepsilon\eta_A = \mu_0/(\rho + d_A + \iota_{A,0})$. Entry then pays whenever the consumption rate exceeds the anchored threshold $\hat{\mu}_A$ of (43). On the investment base the same entry cost is covered when the investment share exceeds its own threshold,

$$\frac{I}{Y} > \hat{v}(d_B) \equiv \mu_0 \frac{\rho + d_B}{\rho + d_A + \iota_{A,0}}, \quad (\text{A.8})$$

which rises with d_B because a shorter-lived rent needs a larger income stream. AI entry stops only when both conditions fail. At the pre-AI baseline the investment share sits below its threshold for any $d_B > 0.10$.

Along the bottleneck path the investment share rises with automation and peaks in the steady state, at $(I/Y)(\phi_\infty)$, which is where the investment base pays most. The boundary

d_B^* is the mortality at which the licence fee just covers the cost at the steady-state investment share, $\hat{v}(d_B^*) = (I/Y)(\phi_\infty)$:

$$d_B^* = (\rho + d_A + \iota_{A,0}) \frac{(I/Y)(\phi_\infty)}{\mu_0} - \rho. \quad (\text{A.9})$$

At the calibration, $d_B^* = 0.13$. Above this threshold, the investment base is never chosen. Below it, once automation has lifted the investment share past the threshold, entrants switch to charging firms, AI entry never stops, automation proceeds to completion, and trend growth ends above the pre-AI rate.

The boundary d_B^* is only about half of d_A because investment is a much smaller stream than consumption. Consumers spend 0.62 of output at the stall, and firms never invest more than 0.35 of output, so a machine licence has to pay for about twice as long as a consumer product to be worth the same entry cost. Measured rents on machines do not last nearly that long. Li and Hall (2020) put R&D depreciation in computers, software and semiconductors at 23 to 36% a year, about two to three times d_B^* , so on the bottleneck path a licence never covers the entry cost and innovators never switch to charging firms.

Appendix B: Robustness

B.1 Perfect foresight


Section 3.3 solves the transition as a temporary equilibrium with static expectations, in which households and entrants assume that the conditions observed at each date will persist. This appendix solves a full perfect-foresight equilibrium, in which all forward-looking decisions are taken with respect to the path the economy actually follows. Capitalists choose consumption from the Euler equation along the path rather than at its balanced-path value, with capital as a state variable and the consumption rate an outcome of goods-market clearing. Entrants price an incumbency as the discounted rents along the path at the equilibrium return on capital, rather than at the balanced-path natural rate of the current state. Free entry holds at each date under both solutions. Workers consume their income each period under both solutions, as in Section 2.3. Perfect foresight leaves the entry threshold and the set of steady states unchanged. Foreseeing the thinning consumer market, capitalists cut consumption on impact, thus shortening the boom and causing AI entry to fade gradually instead of stopping abruptly. The steady state changes by a negligible amount, however.

Agents price the no-AI balanced path before activation and have perfect foresight from activation onward. Using the model equations of Section 2, we solve the whole equilibrium path simultaneously. We do so by a stacked Newton method: the equations at each date of a time grid are stacked into one system in the unknowns of all dates, with the state given at activation and the terminal condition that the economy converges to its steady state, imposed as the balanced-path value of capitalists' wealth at the end of the grid. The procedure converges to the same solution from different starting paths, and the solution is insensitive to the time step and the length of the grid.

Figure B.1 compares the two solutions under the Demand bottleneck scenario. On the news of activation capitalists cut consumption, as they foresee a thinning consumer market and falling rents, and investment rises. The return on capital cannot immediately adjust given the predetermined states, and rises by less than under static expectations to a lower peak due to extra saving. The boom is smaller, with a lower trend growth path that falls below the pre-AI rate sooner, at year 5 rather than year 9. Trend growth reaches a trough after about a decade and recovers slightly as the fall in the consumption rate slows and entry along both dimensions revives marginally. The arrival rate fades as the consumption rate converges to the entry threshold from above. The automated share converges to the lower edge $\bar{\phi}$ of the steady-state set, where the balanced-path consumption rate meets the threshold, marginally below the static steady state.

The same exercise for the Transformative scenario, where AI entry never stops under either solution, shows the same pattern in the boom. Capitalists cut consumption by three tenths on impact, because the consumer market thins faster in that scenario even as output takes off. Trend growth after five decades reaches 5.5 rather than 6.1%, and the return on capital is lower than the static path by about a percentage point. The automation ceiling remains unchanged.

Figure B.1: Static expectations against perfect foresight, Demand bottleneck



Note: The Demand bottleneck path under static expectations and under full perfect foresight. Panel (a): trend growth g_t^* with the natural rate $r_t^* = \rho + \sigma g_t^*$ on the right axis and the pre-AI 2% as a dashed rule. Panel (b): the aggregate AI arrival rate $s_t \iota_{A,t}$. Panel (c): the labour task share $1 - \phi_t$. Panel (d): the consumption rate μ_t (solid) and the AI entry threshold $\hat{\mu}_A$ (dash-dotted), at which entry just pays on a balanced path. Under static expectations AI entry stops where the consumption rate falls through the threshold. Under perfect foresight the arrival rate fades toward zero, and the consumption rate converges to the threshold from above.

B.2 Transformative AI with and without the demand bottleneck

Section 4.2 defines the Transformative scenario as the one that matches most closely the transformative AI in the conventional growth model where capability advances exogenously and independently of the consumer market. This appendix details the construction of such comparisons and their results to highlight the latent role of demand bottleneck in the Transformative scenario.


Consider three scenarios: (i) the Demand bottleneck scenario as in the main text, which uses the central calibration, (ii) transformative AI according to the conventional model where the arrival rate is fixed from activation, obtained under central calibration, and (iii) our model-consistent Transformative scenario, which lets the AI arrival rate respond to consumption rate but with the parameters $(\iota_{A,0}, d_A) = (0.084, 0.001)$ chosen to match the simulated paths of trend growth, labour task share and consumption rate to those of scenario (ii).¹⁹ The difference between (i) and (ii) captures the effects of demand bottleneck at the same parameters. The difference between (ii) and (iii) reflects the latent effect of the demand bottleneck in the Transformative scenario, which holds trend growth back from the accelerating path of the conventional model.

Figure B.2 illustrates the three paths. Under exogenous arrivals the consumption rate falls below the entry threshold of the central calibration within a decade, but because arrivals continue at their activation rate, trend growth keeps accelerating as each arrival raises productivity on a growing share of tasks, reaching 10% after five decades. With the free-entry condition in place, demand bottleneck stops AI entry in the second decade and trend growth settles at 1.6%. Under the same central parameters, the demand bottleneck explains the large and expanding difference between the two paths.

The Transformative scenario is the closest path the model can generate to match the exogenous-arrivals path of scenario (ii). The parameters required for the match are at or beyond the boundaries of the empirical ranges of Section 4.1. AI rents effectively never expire ($d_A = 0.001$, against an empirical range of 0.11–0.49), while the arrival rate of 0.084 implies AI productivity gains about three times those under the Bounded boost scenario. The arrival rate at activation is the rate free entry delivers while the consumption rate is still at its pre-AI level. The conventional transformative path holds the arrival rate fixed at this level even as consumption rate falls, which can be interpreted as optimistic assumptions about rent longevity and AI productivity effects. By design, the model-consistent Transformative scenario and the exogenous-arrivals path are closely aligned, and in the case of trend growth

¹⁹We choose the two parameters $(\iota_{A,0}, d_A)$ to minimise the equally-weighted root mean squared errors of the three variables over the five decades.

Figure B.2: Exogenous arrivals against the Transformative and Demand bottleneck scenarios



Note: The central calibration with the free-entry condition switched off, so that AI innovations arrive at the fixed activation rate of 0.03 a year throughout (teal, dashed), against the Transformative scenario (amber) and the Demand bottleneck scenario (navy) of Section 4.2. Panel (a): trend growth g_t^* with the natural rate $r_t^* = \rho + \sigma g_t^*$ on the right axis, in the model's units. Panel (b): the aggregate AI arrival rate $s_t \iota_{A,t}$. Panel (c): the labour task share $1 - \phi_t$. Panel (d): the consumption rate μ_t (solid) and the AI entry threshold $\hat{\mu}_A$ of each scenario (dash-dotted, in the same colour), lower for Transformative because of its own arrival rate and obsolescence. The exogenous-arrivals path has no entry condition and hence no threshold. All dated in years since activation t_0 . The Transformative pair $(\iota_{A,0}, d_A) = (0.084, 0.001)$ minimises the root mean square distance to the exogenous-arrivals path in panels (a), (c) and (d) over the five decades, with the gaps divided by 5 percentage points, 65 points and 0.5 respectively.

the difference within the first two decades stays within one percentage point.

The remaining differences between the two scenarios reflect the latent effect of demand bottleneck. With endogenous AI entry, the arrival rate rises early but eventually falls. It peaks near 0.06 in the first decade, while the consumption rate is still high, and falls to a third of that after five decades as automation lowers the consumption rate. The Transformative scenario therefore accumulates its capability earlier, and automates about a decade ahead of the exogenous-arrivals path. It entails lower consumption during the early transition.

Trend growth is also faster initially under the Transformative scenario because of the earlier automation. But eventually trend growth tapers off, ending at 6% after five decades, while trend growth continues to accelerate under exogenous arrivals. The reason is precisely the thinning consumer market, which eventually weighs on the AI arrival rate and slows down trend growth even though the arrival rate never falls to zero. This latent effect of the demand bottleneck explains the more moderate growth path in our model-consistent Transformative scenario relative to the exogenous arrival case under the conventional growth model.

B.3 Alternative parameters

Section 4.2 establishes the Demand bottleneck scenario as a robust outcome across the empirical ranges of all five key parameters, $\{d_A, d_H, \iota_{A,0}, \psi, \xi\}$. In this appendix, we further examine the role of each parameter in explaining both the temporary boost and eventual stalling of trend growth under the scenario.

Table B.1 reports the peak in trend growth during the temporary boost, as well as the steady-state trend growth and labour task share, in three columns. We start from the central calibration (first row), change one of the five key parameters at a time to sweep the empirical ranges (second to penultimate rows), and report the ranges of each column under all 32 combinations of the boundary values (last row).

Table B.1: Sensitivity of the Demand bottleneck path to parameters

	Peak trend growth	Steady-state trend growth	Steady-state task share
Central calibration	2.32	1.59	48
$d_A = 0.11$	2.34	1.25	39
$d_A = 0.49$	2.29	1.74	53
$d_H = 0.02$	2.33	1.65	48
$d_H = 0.06$	2.31	1.52	48
$\iota_{A,0} = 0.003$	2.03	1.95	59
$\iota_{A,0} = 0.084$	2.99	1.05	33
$\psi = 4$	2.34	1.60	48
$\psi = 8$	2.29	1.55	47
$\xi = 0.20$	2.28	1.54	47
$\xi = 0.50$	2.35	1.60	49
All combinations	2.02–3.31	0.17–1.98	13–59

Note: Demand bottleneck paths, all values in per cent. Peak is the highest trend growth reached after activation. Steady state is the balanced path at the automated share where AI entry has stopped. Trend growth declines to it monotonically after the peak, so the steady state is also the trough. Task share is the labour task share. The last row gives the lowest and highest value of each column over the 32 combinations of the five parameters at the edges of their ranges. The AI obsolescence edges are the range Li and Hall (2020) measure across industries, the arrival edges invert the lowest and highest published productivity boosts of Section 4.1, the conventional obsolescence edges bracket the Broda and Weinstein (2010) anchor, and the diffusion speed edges are judgment.

The two parameters that raise the early boost are the same ones that deepen the stall. A faster arrival rate at activation, $\iota_{A,0}$ at the upper end of its range, lifts peak trend growth to 3% in the near term but lowers steady-state trend growth to about 1%, with the labour task share settling at a third of all tasks. More durable AI rents, d_A at the lower end, raise peak trend growth by only a couple of basis points but deepen the stall by a third of a point. The difference lies in how the two parameters work. The arrival rate at activation drives AI entry from the start, explaining its large influence on the boom. Rent durability determines how thin the consumer market can get before AI entry stops, hence affecting the duration of AI entry. A later stop implies more automation, a thinner consumer market, smaller rewards for innovation and lower trend growth. Both parameters work through the scale of AI entry, one through its strength and the other through its duration. The other three parameters hardly change the amount of AI entry, and therefore have little effect on the trend growth path. They move the peak and the settled level in the same direction, by a tenth of a point or less.

No combination of the parameters avoids the demand bottleneck pattern. Over all 32 combinations of boundary values, trend growth peaks between 2.0% and 3.3% within a decade of activation and settles between 0.2% and just under 2%, with the steady-state task share between 13% and 59%. The deepest stall combines the fastest arrivals with the most durable rents, the trade-off at its most extreme.

Appendix C: Extensions

C.1 AI lowering research costs

A potential counterforce to the demand bottleneck is that AI could bring capability abundance and lower research costs. If innovating gets cheaper faster than the consumer market thins, the reward on innovation need not fall short of the cost, and AI entry could be sustained indefinitely. The baseline analysis rules this out by pricing research in output, so that the cost of innovating grows with the economy while the prize grows with consumer spending. This appendix examines the case where AI leads to falling research costs, which we model as a decline in the wage paid for researcher time. The wage bill falls relative to output as automation takes over tasks, compressing the cost of innovating. For this argument to hold in our model, the wage of the researchers themselves must fall sharply against output, which is unlikely given the scarcity of AI researchers. Even so, we show that the reward declines more rapidly than the cost unless more than about a third of the cost of AI research is wages.

We conduct two experiments. In the first, a share χ_A of the cost of AI research is priced in wages, while the rest is priced in output. In the second, a share χ_H of the cost of conventional research is priced in wages. Each uses the central calibration with the other share at zero, so the only change from the Demand bottleneck scenario is that part of the research bill moves from being priced in output to being priced in wages.

Research uses a bundle of goods and researcher time with fixed cost shares, so a unit arrival on dimension x costs $\eta_x w^{\chi_x} Y^{1-\chi_x}$ up to a constant. The main text is the case where $\chi_x = 0$, consistent with the lab-equipment form of the growth literature in which research is bought with output, and $\chi_x = 1$ is the knowledge-driven form of Romer (1990) in which research is researchers' time paid at the wage (Rivera-Batiz and Romer, 1991). With one unit of labour the wage bill is $w = \ell Y$, where $\ell \equiv (1 - \phi)(1 - \mu/\varepsilon)$ is the wage-bill share of output, and the cost becomes $\eta_x \ell^{\chi_x} Y$. The prize is unchanged, so free entry is identical to the derivation of (29) except that the consumption rate μ is replaced by μ/ℓ^{χ_x} , the consumption rate per unit of the wage-bill share, which we call the adjusted consumption rate. Anchoring at the pre-AI balanced path as in Section 4.1 gives

$$\iota_x = \max \left\{ 0, (\rho + d_x + \iota_{x,0}) \frac{\mu}{\mu_0} \left(\frac{\ell_0}{\ell} \right)^{\chi_x} - (\rho + d_x) \right\}. \quad (\text{C.1})$$

Automation now pulls the arrival rate in both directions. It lowers the consumption rate, which weakens entry as before, and it lowers the wage-bill share, which cheapens wage-priced research and strengthens entry. Which direction dominates depends on χ_x .

A boundary condition for χ_A can be written in closed form. The wage bill falls as tasks are automated, so research priced in wages is cheapest at full automation, where the consumer market is also thinnest. At this point, AI entry does not pay if the consumption rate has fallen further than the wage-priced part of the cost, a condition that can be written in terms of the share of the research cost:

$$\chi_A < \chi_A^* \equiv \frac{\ln(\mu_0/\mu(\phi_{\max})) - \ln((\rho + d_A + \iota_{A,0})/(\rho + d_A))}{\ln(\ell_0/\ell(\phi_{\max}))}. \quad (\text{C.2})$$

At the calibration, $\chi_A^* = 0.22$. Below this threshold, entry is not worthwhile at full automation where the cost is cheapest, so it must stop at some earlier point, and the stall is certain. Above the threshold, entry pays at full automation, but the economy may nonetheless never get there. Early in the transition, the consumption rate falls faster than the wage bill, so the adjusted consumption rate falls and entry weakens. Once entry stops, automation can go no further than the frontier already reached, short of the shares at which entry would be worthwhile again. Whether entry survives long enough to reach full automation requires simulation. We report this in Table C.1.

Table C.1: Research costs priced in wages

	AI entry	Steady-state trend growth	Steady-state task share
Share of AI research priced in wages, χ_A			
0 (Demand bottleneck)	stops	1.59	48
0.1	stops	1.51	46
0.2	stops	1.37	42
0.3	stops	0.99	31
0.4	continues	4.98	5
0.5	continues	9.00	5
Share of conventional research priced in wages, χ_H			
0.2	stops	1.72	49
0.5	stops	1.90	49
0.7	stops	2.02	50

Note: Central calibration with the other share at zero. AI entry stops when the arrival rate dies away, and continues when it stays positive and automation reaches its ceiling. Steady state is the balanced path at the automated share where entry has died away or, where it continues, at the automation ceiling. Trend growth and the labour task share are in per cent. The simulated boundary above which entry continues is about 0.31, against the closed-form 0.22.

The table shows that two polar outcomes obtain. When less than a third of the AI research cost is priced in wages ($\chi_A \leq 0.3$), AI entry stops, and the economy settles at trend growth below the pre-AI rate, as in the Demand bottleneck scenario. With the share of AI research priced in wages higher than a third, AI entry never stops, automation runs to its ceiling, and trend growth ends at 5 to 9% with the labour task share falling to 5%. These

two outcomes are the two polar ends of a race. Automation thins the consumer market, which weakens the incentive to enter, and it lowers the wage bill, which cheapens wage-priced research. Early in the transition the consumer market thins faster than the wage bill falls, so entry weakens regardless of the share. Near full automation, however, the wage bill falls faster. Cheaper research can therefore rescue AI entry only if entry survives long enough to reach sufficiently high automation. And with less than a third of the cost in wages, this does not materialise.

Within the range $\chi_A \leq 0.3$, cheaper research in fact makes the outcome worse. It keeps AI entry alive for longer, so more tasks are automated before entry stops, but the wage bill never falls fast enough for growth to take off. The economy then settles with fewer tasks left for labour and a thinner consumer market, leaving conventional innovation earning less. Steady-state trend growth falls from 1.6% when no research is priced in wages to 1% when $\chi_A = 0.3$, as the labour task share falls from 48 to 31%.

The share of conventional research priced in wages, χ_H , has little bearing on AI entry, which stops in all cases considered in the table. Once AI entry has stalled, the automated share and the wage bill stop moving, so a cheaper conventional research bill lifts conventional entry rather than racing a thinning market. Trend growth is therefore increasing in χ_H , and exceeds the pre-AI rate only when more than two thirds of the cost of conventional research is priced in wages. Meanwhile, the labour task share hardly moves because automation stops where it did.

C.2 Automated research

The main text holds the cost of AI research fixed as capability grows. In this appendix, we let each advance in AI capability lower the cost of the next improvement, and ask whether this self-improvement can help AI entry survive the thinning consumer market. The answer is that each doubling of capability must cut the cost of the next advance by well over half. Short of that, cheaper research sustains AI entry for longer but does not save it from the thinning consumer market, and the economy ends up worse off than with no self-improvement at all.

We let the entry cost on the AI dimension fall with capability,

$$\eta_A(A_t) = \eta_{A,0} (A_t/A_0)^{-\gamma}, \quad (\text{C.3})$$

where $\gamma \geq 0$ indexes the strength of self-improvement and $\gamma = 0$ is the case in the main text. The consumption rate below which AI entry stops, $\hat{\mu}_A$, is proportional to the entry cost by

(30), so it falls at the same rate, $\hat{\mu}_A(A_t) = \hat{\mu}_{A,0} (A_t/A_0)^{-\gamma}$, and AI entry survives a thinner consumer market the more capability has accumulated.

Whether AI entry survives is again a race, this time between capability and the consumer market. AI entry is safe once capability has grown enough for the threshold to fall below the thinnest consumer market the economy can reach, namely the consumption rate at full automation $\mu(\phi_{\max})$:

$$\ln \frac{A_t}{A_0} > X_{\text{crit}}(\gamma) \equiv \frac{1}{\gamma} \ln \frac{\hat{\mu}_{A,0}}{\mu(\phi_{\max})}. \quad (\text{C.4})$$

If AI entry stops before capability reaches this point, it stops for good, because a stop freezes capability and with it the threshold, while adoption keeps thinning the consumer market. The critical degree of self-improvement γ^* is the one at which capability reaches the point in (C.4) exactly when entry would otherwise stop. At the calibration $\gamma^* = 1.27$. Below it AI entry stops and trend growth ends below the pre-AI rate. Above it AI entry never stops and trend growth diverges. Figure C.1 draws both cases, at $\gamma = 0$ and 1 on one side of the critical degree and $\gamma = 1.5$ on the other.



When self-improvement is weaker than the critical degree, it makes the outcome worse, as a wage-priced share below a third did in Appendix C.1. Cheaper research keeps AI entry alive for longer, so more tasks are automated before entry stops, and the economy settles with fewer tasks left for labour and a thinner consumer market. Steady-state trend growth falls from 1.6% with no self-improvement to 1.3% at $\gamma = 1$ and to 0.9% at the critical degree itself, and the labour task share falls from 48 to 28%.

When self-improvement is stronger than the critical degree, AI entry survives, but for decades the path is indistinguishable from a bottleneck. At $\gamma = 1.5$ the AI arrival rate never reaches zero, but it falls for decades, because early in the transition the consumer market thins faster than the threshold falls, and trend growth falls with it, to below the level at which the Demand bottleneck scenario settles. Capability keeps accumulating through this period, so the threshold keeps falling, and it eventually falls faster than the consumer market thins. The arrival rate then turns up, trend growth climbs back through the pre-AI rate, and it rises without limit thereafter, with almost nine tasks in ten automated by the end of the figure's horizon.


The critical degree is a statement about technology. For AI entry to survive, each doubling of capability must cut the cost of the next advance by a factor of about $2^{\gamma^*} \approx 2.4$, and this must hold at every level of capability, without end. Research productivity has instead been falling for decades (Bloom et al., 2020).

Figure C.1: Automated research: three degrees of self-improvement


(a) Trend growth (g^*) and the natural rate (r^*) (b) AI arrival rate (ι_A)



(c) Labour task share ($1 - \phi$)



(d) Consumption rate (μ) and $\hat{\mu}_A(t)$



— Demand bottleneck, $\gamma=0$ — Automated research, $\gamma=1.5$
— Automated research, $\gamma=1.0$ - - - AI entry threshold $\hat{\mu}_A(t)$

Note: The Demand bottleneck path ($\gamma = 0$) and the automated-research paths at $\gamma = 1$ and $\gamma = 1.5$ over 119 years since activation t_0 . Panel (a): trend growth g_t^* with the natural rate $r_t^* = \rho + \sigma g_t^*$ on the right axis, in the model's units, with the pre-AI 2% as a dashed rule. Panel (b): the AI arrival rate $s_t \iota_{A,t}$. Panel (c): the labour task share $1 - \phi_t$. Panel (d): the consumption rate μ_t (solid) and the AI entry threshold $\hat{\mu}_A(A_t)$ (dash-dotted) in the same colour. AI entry stops where the two meet.

References

- ACEMOGLU, D. (2024): “The simple macroeconomics of AI,” NBER Working Paper 32487, National Bureau of Economic Research.
- ACEMOGLU, D. AND S. JOHNSON (2023): *Power and progress: Our thousand-year struggle over technology and prosperity*, New York: PublicAffairs.
- ACEMOGLU, D. AND J. LINN (2004): “Market size in innovation: Theory and evidence from the pharmaceutical industry,” *Quarterly Journal of Economics*, 119, 1049–1090.
- ACEMOGLU, D. AND P. RESTREPO (2018): “The race between man and machine: Implications of technology for growth, factor shares, and employment,” *American Economic Review*, 108, 1488–1542.
- (2019): “Automation and new tasks: How technology displaces and reinstates labor,” *Journal of Economic Perspectives*, 33, 3–30.
- (2022): “Tasks, automation, and the rise in U.S. wage inequality,” *Econometrica*, 90, 1973–2016.
- AGHION, P. AND P. HOWITT (1992): “A model of growth through creative destruction,” *Econometrica*, 60, 323–351.
- AGHION, P., B. JONES, AND C. JONES (2019): “Artificial intelligence and economic growth,” in *The economics of artificial intelligence: An agenda*, ed. by A. Agrawal, J. Gans, and A. Goldfarb, University of Chicago Press, 237–282.
- ALDASORO, I., S. DOERR, L. GAMBACORTA, AND D. REES (2024): “The impact of artificial intelligence on output and inflation,” BIS Working Paper 1179, Bank for International Settlements.
- ANDREWS, I. AND M. FARBOODI (2025): “Do markets believe in transformative AI?” NBER Working Paper 34243, National Bureau of Economic Research.
- AUTOR, D., C. CHIN, A. SALOMONS, AND B. SEEGMILLER (2024): “New frontiers: The origins and content of new work, 1940–2018,” *Quarterly Journal of Economics*, 139, 1399–1465.
- BANERJEE, S. AND J. WANG (2026): “The reverse big push: Generative AI and self-fulfilling automation,” ArXiv:2608.25602.

- BAQAEE, D. AND E. FARHI (2020): “Productivity and misallocation in general equilibrium,” *Quarterly Journal of Economics*, 135, 105–163.
- BARRO, R. AND X. SALA-I-MARTIN (2004): *Economic growth*, MIT Press, 2 ed.
- BAYRAKTAR, E. (2026): “Automation, income incidence, and capital accumulation in incomplete markets,” ArXiv:2605.05127; earlier circulated as “The demand externality for automation”.
- BENIGNO, G. AND L. FORNARO (2018): “Stagnation traps,” *Review of Economic Studies*, 85, 1425–1470.
- BICK, A., A. BLANDIN, AND D. J. DEMING (2026): “The rapid adoption of generative AI,” *Management Science*, 72, 8440–8456.
- BLOOM, N., C. JONES, J. VAN REENEN, AND M. WEBB (2020): “Are ideas getting harder to find?” *American Economic Review*, 110, 1104–1144.
- BRIGGS, J. AND D. KODNANI (2023): “The potentially large effects of artificial intelligence on economic growth,” Global economics analyst, Goldman Sachs.
- BRODA, C. AND D. WEINSTEIN (2010): “Product creation and destruction: Evidence and price implications,” *American Economic Review*, 100, 691–723.
- BRYNJOLFSSON, E. (2022): “The Turing trap: The promise and peril of human-like artificial intelligence,” *Daedalus*, 151, 272–287.
- BRYNJOLFSSON, E., D. LI, AND L. RAYMOND (2025): “Generative AI at work,” *Quarterly Journal of Economics*, 140, 889–942.
- BUREAU OF ECONOMIC ANALYSIS (2026): “Table 1.1.10. Percentage shares of gross domestic product,” National Income and Product Accounts, U.S. Department of Commerce.
- CHOW, T., B. HALPERIN, AND J. Z. MAZLISH (2026): “Transformative AI, existential risk, and real interest rates,” Working paper; first posted January 2023.
- CHRISTENSEN, J. H. E. AND G. D. RUDEBUSCH (2026): “Can fiscal, AI, or monetary news explain the rise in r^* ?” Working Paper 2026-19, Federal Reserve Bank of San Francisco.
- COMIN, D. AND B. HOBIJN (2010): “An exploration of technology diffusion,” *American Economic Review*, 100, 2031–2059.

- CORRADO, C., J. HASKEL, C. JONA-LASINIO, AND M. IOMMI (2022): “Intangible capital and modern economies,” *Journal of Economic Perspectives*, 36, 3–28.
- CORRADO, C., C. HULTEN, AND D. SICHEL (2009): “Intangible capital and U.S. economic growth,” *Review of Income and Wealth*, 55, 661–685.
- DAVIDSON, T., B. HALPERIN, T. HOULDEN, AND A. KORINEK (2026): “When does automating AI research produce explosive growth? Feedback loops in innovation networks,” NBER Working Paper 35155, National Bureau of Economic Research.
- DEL NEGRO, M., D. GIANNONE, M. GIANNONI, AND A. TAMBALOTTI (2017): “Safety, liquidity, and the natural rate of interest,” *Brookings Papers on Economic Activity*, 235–316.
- DELL’ACQUA, F., E. MCFOWLAND III, E. R. MOLLICK, H. LIFSHITZ-ASSAF, K. C. KELLOGG, S. RAJENDRAN, L. KRAYER, F. CANDELON, AND K. R. LAKHANI (2023): “Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality,” Working Paper 24-013, Harvard Business School.
- EDMOND, C., V. MIDRIGAN, AND D. XU (2023): “How costly are markups?” *Journal of Political Economy*, 131, 1619–1675.
- ELOUNDOU, T., S. MANNING, P. MISHKIN, AND D. ROCK (2024): “GPTs are GPTs: Labor market impact potential of LLMs,” *Science*, 384, 1306–1308.
- ERDIL, E. AND T. BESIROGLU (2023): “Explosive growth from AI automation: A review of the arguments,” ArXiv:2309.11690.
- FERNALD, J. AND C. JONES (2014): “The future of US economic growth,” *American Economic Review*, 104, 44–49.
- FILIPPUCCI, F., P. GAL, AND M. SCHIEF (2024): “Miracle or myth? Assessing the macroeconomic productivity gains from artificial intelligence,” Artificial Intelligence Papers 29, OECD.
- FOELLM, R. AND J. ZWEIMÜLLER (2006): “Income distribution and demand-induced innovations,” *Review of Economic Studies*, 73, 941–960.
- FORNARO, L. AND M. WOLF (2022): “Monetary policy in the age of automation,” Working paper, CREI, September 2022.
- (2026): “Macroeconomic policies for AI,” Discussion Paper 21412, CEPR.

- GARCIA-MACIA, D., C.-T. HSIEH, AND P. KLENOW (2019): “How destructive is innovation?” *Econometrica*, 87, 1507–1541.
- GREENWOOD, J., Z. HERCOWITZ, AND P. KRUSELL (1997): “Long-run implications of investment-specific technological change,” *American Economic Review*, 87, 342–362.
- GRIES, T. AND W. NAUDÉ (2021): “The race of man and machine: Implications of technology when abilities and demand constraints matter,” IZA Discussion Paper 14341, IZA Institute of Labor Economics.
- GROSSMAN, G. AND E. HELPMAN (1991): *Innovation and growth in the global economy*, MIT Press.
- HALL, R. (1988): “The relation between price and marginal cost in U.S. industry,” *Journal of Political Economy*, 96, 921–947.
- HÉMOUS, D. AND M. OLSEN (2022): “The rise of the machines: Automation, horizontal innovation, and income inequality,” *American Economic Journal: Macroeconomics*, 14, 179–223.
- HOLSTON, K., T. LAUBACH, AND J. WILLIAMS (2023): “Measuring the natural rate of interest after COVID-19,” Staff Report 1063, Federal Reserve Bank of New York.
- HOWITT, P. (1999): “Steady endogenous growth with population and R. & D. inputs growing,” *Journal of Political Economy*, 107, 715–730.
- HUMLUM, A. AND E. VESTERGAARD (2025): “Still waters, rapid currents: Early labor market transformation under generative AI,” NBER Working Paper 33777, National Bureau of Economic Research.
- JARAVEL, X. (2019): “The unequal gains from product innovations: Evidence from the U.S. retail sector,” *Quarterly Journal of Economics*, 134, 715–783.
- JONES, C. (1995): “R&D-based models of economic growth,” *Journal of Political Economy*, 103, 759–784.
- JONES, C. AND C. TONETTI (2026): “Past automation and future AI: How weak links tame the growth explosion,” Working paper, Stanford Graduate School of Business.
- JONES, C. I. (2024): “The A.I. dilemma: Growth versus existential risk,” *American Economic Review: Insights*, 6, 575–590.

- (2026): “A.I. and our economic future,” *Journal of Economic Perspectives*, 40, 3–22, also NBER Working Paper 34779.
- KORINEK, A., C. I. JONES, S. SACHER, T. COTTER, AND P. MCCRORY (2026): “Economic scenarios for transformative AI,” Working Paper 2026-02, The Anthropic Institute.
- KORINEK, A. AND D. SUH (2024): “Scenarios for the transition to AGI,” NBER Working Paper 32255, National Bureau of Economic Research.
- KUNG, H., H. LUSTIG, AND J. PARON (2026): “U.S. Treasury investors are long in AI,” Working paper.
- LAUBACH, T. AND J. WILLIAMS (2003): “Measuring the natural rate of interest,” *Review of Economics and Statistics*, 85, 1063–1070.
- LEE, B. (2023): “Wealth inequality and endogenous growth,” *Journal of Monetary Economics*, 133, 132–148.
- LI, W. AND B. HALL (2020): “Depreciation of business R&D capital,” *Review of Income and Wealth*, 66, 161–180.
- MARESCA, C. (2026): “AGI could lower interest rates,” Working paper.
- MATSUYAMA, K. (2002): “The rise of mass consumption societies,” *Journal of Political Economy*, 110, 1035–1070.
- MELINA, G. AND S. VILLA (2025): “From servers to rates: AI, ICT capital, and the natural rate,” IMF Working Paper 2025/224, International Monetary Fund.
- MIAN, A., L. STRAUB, AND A. SUFI (2021): “Indebted demand,” *Quarterly Journal of Economics*, 136, 2243–2307.
- MOLL, B., L. RACHEL, AND P. RESTREPO (2022): “Uneven growth: Automation’s impact on income and wealth inequality,” *Econometrica*, 90, 2645–2683.
- MURPHY, K., A. SHLEIFER, AND R. VISHNY (1989): “Income distribution, market size, and industrialization,” *Quarterly Journal of Economics*, 104, 537–564.
- NATIONAL CENTER FOR SCIENCE AND ENGINEERING STATISTICS (2026): “National patterns of R&D resources: 2023–24 data update,” NSF 26-313, National Science Foundation, Alexandria, VA.

- NORDHAUS, W. D. (2021): “Are we approaching an economic singularity? Information technology and the future of economic growth,” *American Economic Journal: Macroeconomics*, 13, 299–332.
- NOY, S. AND W. ZHANG (2023): “Experimental evidence on the productivity effects of generative artificial intelligence,” *Science*, 381, 187–192.
- RIVERA-BATIZ, L. AND P. ROMER (1991): “Economic integration and endogenous growth,” *Quarterly Journal of Economics*, 106, 531–555.
- ROMER, P. (1990): “Endogenous technological change,” *Journal of Political Economy*, 98, S71–S102.
- SCHMOOKLER, J. (1966): *Invention and economic growth*, Harvard University Press.
- TRAMMELL, P. AND A. KORINEK (2026): “Economic growth under transformative AI,” *Annual Review of Economics*, 18, 589–611.
- YOUNG, A. (1998): “Growth without scale effects,” *Journal of Political Economy*, 106, 41–63.
- ZEIRA, J. (1998): “Workers, machines, and economic growth,” *Quarterly Journal of Economics*, 113, 1091–1117.