Catch-up growth after AI

Technical note for the calculator

Rafael Latham-Proença

August 2026

This note documents the machinery behind the calculator that accompanies the blog post "The staggering welfare consequences of catch-up growth after AI." The note has two parts. First, I use a simple model of automation to show how a single driving force, the quality-adjusted cost of automation falling over time, gives rise to a growing growth rate and a fixed lag between frontier and non-frontier economies. Then I explain how I calibrate the model to generate the central results in the calculator introduced in the post.

1 Where the growth acceleration comes from

1.1 A task economy

Production is a long list of necessary tasks, and machines gradually become able to do more of them. The basic idea goes back at least to Zeira (1998). Acemoglu and Autor (2011), Acemoglu and Restrepo (2018), and Aghion, Jones, and Jones (2019) develop it in detail.

Index tasks by $i \in [0, 1]$. Think of driving a truck, welding a joint, or filing taxes. For each task, people and machines are alternative ways to produce the same output. Let $A_j^L(i)$ and $A_j^K(i,t)$ be their productivities in country j, let w_j be the wage, and let R_j be the annualized rental price of machine capital. Unit costs are

$$c_j^L(i) = \frac{w_j}{A_j^L(i)}, \qquad c_j^K(i,t) = \frac{R_j}{A_j^K(i,t)}.$$

A cost-minimizing firm assigns task i to machines when

$$i \in \mathcal{M}_j(t) \quad \Longleftrightarrow \quad c_j^K(i,t) \le c_j^L(i).$$

This is the standard assignment rule in task models. Order tasks by machines' comparative advantage. There is then a threshold $f_j(t)$ such that machines perform tasks below the threshold and people perform those above it. Because the task interval has length one, $f_j(t)$ is also the fraction of tasks performed by machines. It is the same $f(t)$ used in the blog.

Suppress country and time subscripts for the moment. Let $y(i)$ be the amount of task i that gets done. Aggregate output is a Cobb–Douglas aggregate over the continuum of tasks:

$$Y = \exp\left(\int_0^1 \ln y(i)\, di\right). \tag{1}$$

The tasks are complements. Doing a great deal of one cannot make up for skipping another.

Machines occupy $i \in [0, f]$, and people do the rest. Normalize so that one unit of capital produces one unit of a machine task and one unit of effective labor produces one unit of a human task:

$$
\begin{gathered}
y(i)=\begin{cases} k(i), & 0 \le i \le f, \\ A\ell(i), & f < i \le 1, \end{cases} \\
\int_0^f k(i)\,di = K, \qquad \int_f^1 \ell(i)\,di = L.
\end{gathered} \tag{2}
$$

Here K is the capital stock, L is the number of workers, and A is labor productivity. The aggregator is concave in each task input, so output is maximized by spreading each input evenly across its tasks. Thus $k(i) = K/f$ and $\ell(i) = L/(1-f)$. If machines perform 40% of tasks, the capital stock is spread across that 40%, while labor is spread across the remaining 60%. Substituting into the aggregator gives

$$
\ln Y = f \ln\left(\frac{K}{f}\right) + (1-f)\ln\left(\frac{AL}{1-f}\right), \tag{3}
$$

which exponentiates to

$$
Y = \left(\frac{K}{f}\right)^{f}\left(\frac{AL}{1-f}\right)^{1-f}. \tag{4}
$$

The exponents are the machine and human shares of tasks. The edge cases $f = 0$ and $f = 1$ are understood as limits.

Capital accumulates through investment, $\dot{K} = sY - \delta K$, where s is the investment share and δ is depreciation. While people still perform some necessary tasks, accumulated machines run into a human bottleneck and returns to capital diminish. As $f \to 1$, that bottleneck disappears.

To describe the endpoint, let q_K denote the annual gross output produced by one unit of reproducible machine capital. In the absence of a binding human or other fixed-factor constraint, full automation gives

$$
Y = q_K K.
$$

If the investment share is fixed, capital and output then grow at

$$
g_{\text{full}} = s q_K - \delta.
$$

Conceptually, q_K is the annual gross output produced by one unit of reproducible machine capital. It is not itself the capital-reproduction rate: sq_K is the rate at which output is converted into new capital, and $sq_K - \delta$ is the resulting net growth rate. Its inverse, $1/q_K$, is the capital-output ratio: for example, $q_K = 0.5$ means that producing one dollar of annual output requires two dollars of machine capital. This AK-style feedback is what makes a large takeoff possible in principle. It does not, by itself, determine how large the takeoff is: that also depends on investment behavior, depreciation, and whether land, energy, materials, or other irreproducible factors become binding (Trammell and Korinek, 2026).[1]

1.2 Falling machine costs create the takeoff and the lag

The assignment rule above says that automation pays when a machine does a task more cheaply than a person. The rest of the model keeps track of when each task crosses that line.

[1]Equations (2)–(4) normalize the productivity of machine capital to one for notational convenience. I reintroduce q_K here because its scale matters for growth at the full-automation endpoint.

On the labor side, I simplify task-specific labor costs to a country-level unit labor cost, $c^L_j = w_j/A^L_j$, the wage divided by worker productivity.[2] On the machine side, let t be years since the frontier automation wave begins. Specialize the machine unit cost introduced above to

$$c^K_j(i,t) = \omega_j \kappa_i e^{-\lambda t}, \tag{5}$$

where κ_i measures how hard the task is to automate and λ is the annual rate at which quality-adjusted machine costs fall. The country wedge ω_j covers anything that makes machines expensive or ineffective locally. It includes financing costs, scarce foreign exchange, tariffs, unreliable power, missing engineers, and incompatible capital or data.[3] The exponential specification is a description of the automation transition, not a claim that the sticker price of physical machines literally falls to zero forever. The calculator truncates the process at T and summarizes the effective endpoint through $f_{\max}$.

A cost-minimizing firm automates task i once $c^K_j(i,t) \le c^L_j$. Equivalently,

$$\kappa_i \le \frac{c^L_j}{\omega_j} e^{\lambda t}. \tag{6}$$

As machine costs fall, this profitability threshold sweeps through increasingly difficult tasks. If F_κ is the distribution of task difficulty, the machine-task share is

$$f_j(t) = F_\kappa\left(\frac{c^L_j}{\omega_j} e^{\lambda t}\right).$$

It traces an S-curve when log task difficulty is distributed roughly logistically (Appendix A.2). The calculator does not ask users to choose λ or a difficulty distribution. It summarizes the frontier transition with a visible takeoff length T, an initial machine-task share f_0, and an automation ceiling $f_{\max}$. For $0 < t < T$, the path is

$$f_F(t) = f_0 + (1 - f_0) f_{\max} \frac{\sigma(6(t/T - 1/2)) - \sigma(-3)}{\sigma(3) - \sigma(-3)}, \qquad \sigma(u) = \frac{1}{1 + e^{-u}}. \tag{7}$$

Before the window $f_F(t) = f_0$, and after it the task share remains at its endpoint.[4]

Now compare the frontier (F) with a poorer economy (P). Suppose they face the same task-difficulty distribution and the same global λ, and differ only in labor costs and the deployment wedge. Their automation paths then have the same shape, with the poorer economy shifted right by D years:[5]

$$f_P(t) = f_F(t - D), \qquad D = \frac{1}{\lambda} \ln\left(\frac{\omega_P / c^L_P}{\omega_F / c^L_F}\right). \tag{8}$$

The lag grows when machines are costlier to deploy in the poorer economy, when its labor is cheaper per unit of task output, or when machine costs improve more slowly.

[2] Unit labor cost is the relevant comparison. A Malawian accountant may earn around one-thirtieth of a Seattle accountant's wage but also process fewer accounts per year. The cost per account differs by much less than 30 to 1, so unit-labor-cost gaps across countries are much smaller than raw wage gaps.

[3] A missing complementary input can be written as a higher cost or, equivalently, as lower machine productivity.

[4] Truncating and rescaling the logistic tails is only a convenience. The economic mechanism comes from the cost threshold. The factor 6 is a rounded version of the 5–95% transition width derived in Appendix A.2. The blog motivates the same S-curve with contagion logic, where the pace of automation is proportional to $f(1 - f)$. The logistic also solves that differential equation, so the two readings coincide.

[5] The letter D is borrowed from Comin and Mestieri (2018). Their estimated adoption lags are the closest empirical cousins of this parameter, and Section 2.5 maps one to the other.

The blog lists four reasons for delayed adoption. Three fit directly into these country terms. Expensive machines, scarce dollars, and missing complementary inputs raise ω_P. Low wages lower c_P^L. The remaining reason, the more physical composition of poor-country output, changes the task mix and sits outside this threshold model.[6]

If machine costs halve every h years, then $\lambda = \ln 2/h$ and

$$D = h \log_2 \left(\frac{\omega_P / c_P^L}{\omega_F / c_F^L} \right). \tag{9}$$

It is useful to read equation (9) backwards. The central calibration of $D = 10$ years, combined with a three-year halving time for machine costs, implies that the poorer economy faces a machine-to-labor cost ratio about $2^{10/3} \approx 10$ times less favorable than the frontier's.

1.3 From tasks to the growth path

Let f_0 be the fraction of tasks performed by machines when the modeled automation wave begins. Let $f_{\max} \in [0, 1]$ be the fraction of the remaining human task mass that eventually enters the wave. The endpoint is therefore

$$f_\infty = f_0 + (1 - f_0) f_{\max}. \tag{10}$$

The same $f_j(t)$ now does all the work. It is the task boundary in the assignment problem, the fraction of tasks performed by machines, and the exponent on capital in equation (4). The calculator does not need to estimate f_0 because observed baseline growth already reflects whatever automation exists when the modeled wave begins.

Why calibrate the path for f directly? Equation (5) suggests a more structural route. I could choose λ, specify the distribution of task difficulty, and let falling machine costs generate $f(t)$. I do not think those inputs can be measured well enough for this exercise. Quality-adjusted machine costs are hard to compare across software, robots, and the complementary capital needed to use them. The task-difficulty distribution is harder still. So the calculator calibrates $f(t)$ through T and $f_{\max}$. I use the falling-cost model to motivate the S-shape and to discipline the lag D.

The calculator maps frontier automation into growth with a geometric interpolation between two pinned endpoints. Growth before the takeoff is g_0 at $f_F = f_0$. The full-automation benchmark is g_{full} at $f_F = 1$:

$$g_F(t) + \delta = (g_0 + \delta)^{1 - \frac{f_F(t) - f_0}{1 - f_0}} (g_{\text{full}} + \delta)^{\frac{f_F(t) - f_0}{1 - f_0}}. \tag{11}$$

I interpolate gross growth, $g + \delta$, because the task model describes gross capital reproduction. Appendix A.1 motivates this interpolation from the capital-accumulation equation and states the additional approximation required.[7] Equation (11) makes gross growth log-linear in the fraction of the initially human task mass that has been automated. Equal increments of $(f_F - f_0)/(1 - f_0)$ multiply $g_F + \delta$ by equal factors. Under the central calibration, each ten-percentage-point increment multiplies gross growth by $e^{\beta/10} \approx 1.13$, where $\beta = \ln[(g_{\text{full}} + \delta)/(g_0 + \delta)]$.

There is no separate assumption about the time shape of $g_F(t)$. Once the S-shaped path for $f_F(t)$ and the endpoint mapping in equation (11) are fixed, the growth path follows.

The endpoint reached inside the window is

$$g_{\text{end}} + \delta = (g_0 + \delta)^{1 - f_{\max}} (g_{\text{full}} + \delta)^{f_{\max}}. \tag{12}$$

[6] The derivation assumes both economies face the same task mix, automation ceiling, and λ. When those differ, a single constant D will not fit the whole transition, so the calculator's D is best read as an effective average lag. Equation (8) also prices profitability alone. The empirical lag adds installation, reorganization, and intensity of use, which Section 2.5 discusses.

[7] Two features follow from the pinning. The initial task share f_0 never needs a number because observed g_0 already contains it. The mapping also needs no separate curvature dial. The two endpoints determine it.

With $g_0 = 2\%$, $g_{\text{full}} = 18\%$, and $\delta = 5\%$, setting $f_{\max} = 0.75$ gives $g_{\text{end}} = 12.1\%$. Setting it to 0.50 gives 7.7%. The full benchmark applies only when $f_{\max} = 1$.

Finally, the lag carries through unchanged. Equation (8) says that the poorer economy follows the same machine-task share D years later. Applying the same automation-to-growth mapping gives it the same growth acceleration D years later.

1.4 The gap created by delayed adoption

The point of the model above is to derive the two assumptions stated at the beginning of the blog post and then show what follows from them.

The first assumption is that AI raises frontier growth during the takeoff. Equation (7) gives an increasing frontier machine-task share $f_F(t)$. Equation (11) maps a higher machine-task share into higher growth whenever $g_{\text{full}} > g_0$. Frontier growth therefore rises from g_0 to g_{end}. Define the resulting AI growth boost as

$$b(t) \equiv g_F(t) - g_0, \tag{13}$$

with $b(t) = 0$ before the takeoff starts. During the takeoff, $b(t)$ is positive and increasing. This is the first assumption in the blog.

The second assumption is that developing countries eventually receive the same growth boost, but with a lag. Equation (8) gives $f_P(t) = f_F(t - D)$. If the same machine-task share produces the same AI growth boost in both economies, the developing economy receives $b(t - D)$ when the frontier receives $b(t)$. This is the second assumption in the blog. The developing economy can still have its own ordinary growth rate. Let $\bar{g}_P$ denote its growth rate in the absence of the AI boost.

Those two assumptions are sufficient to generate an expanding delay-induced income gap during the takeoff. To isolate that effect, compare two paths for the same developing economy. Both start from income y_0 and share the baseline $\bar{g}_P$. The path $y^0(t)$ receives the AI boost without delay, while $y^D(t)$ receives the same boost D years later:

$$\frac{d \ln y^0(t)}{dt} = \bar{g}_P + b(t), \tag{14}$$

$$\frac{d \ln y^D(t)}{dt} = \bar{g}_P + b(t - D). \tag{15}$$

Integrating and subtracting gives

$$\ln\left(\frac{y^0(t)}{y^D(t)}\right) = \int_{t-D}^{t} b(s)\, ds, \tag{16}$$

where $b(s) = 0$ for $s < 0$. The log income gap is the part of the AI boost that the immediate-adoption path has already received and the delayed path has not. Starting income and ordinary developing-country growth cancel because both paths share them.[8]

Differentiating makes the expansion result explicit:

$$\frac{d}{dt} \ln\left(\frac{y^0(t)}{y^D(t)}\right) = b(t) - b(t - D).$$

[8] This result isolates the gap caused by delayed adoption. The observed gap between a frontier economy and a developing economy also reflects any difference between their baseline growth rates. The delay component still evolves exactly as shown here.

This expression is positive throughout the detachment window. At first, the immediate-adoption path receives a positive boost while the delayed path receives none. Once both are in the takeoff, the immediate path is always D years further along an increasing boost curve. After the immediate path reaches the endpoint at T, it receives the full boost while the delayed path is still catching up. The gap therefore keeps expanding until year $T + D$.

Let $B = g_{\text{end}} - g_0$ be the boost once the takeoff is complete. After year $T + D$, both paths receive B, so their growth rates are again equal and the gap stops widening at

$$\frac{y^0}{y^D} = e^{BD}. \tag{17}$$

Under the central settings, $B = 16$ percentage points and $D = 10$ years, so $e^{1.6} = 4.95$. By year 25, income in the immediate-adoption path is almost five times income in the delayed-adoption path. The model then keeps the gap fixed. It contains no force that closes it.

2 Calibration

I chose central values that are easy to inspect and argue with. I do not pretend the five main inputs have the same empirical basis. Table 1 says what supports each one.

Table 1: The main scenario inputs

Parameter	Central	Sensitivity shown	What supports it?
T	15 years	5–30	A reduced-form summary of machine-cost decline and task-difficulty dispersion, informed by AI progress and slower economy-wide deployment.
$f_{\max}$	100%	50–100%	No useful point estimate. A visible robustness parameter for the effective automation reached by the end of the modeled transition.
g_{full}	18%	10–30%	The steady-state growth rate in Korinek and Suh's numerical full-automation scenarios; imported as a reduced-form benchmark rather than independently estimated here.
D	10 years	5–20	Frontier-relative arrival differentials of 6–8 years for recent technologies, plus a widening intensity-of-use margin (Comin and Mestieri, 2018). Equation (8) provides a separate cross-check.
H	25 years	User choice	The accounting window. The central value is $T + D$.

2.1 Baseline growth rates, g_0 and $\bar{g}_P$

The central settings are $g_0 = 2.0\%$ and $\bar{g}_P = 2.5\%$ per year. These are rounded working values, and the two groups can of course grow at very different rates.[9] The 2.5% input covers low- and middle-income economies excluding China. An Africa-specific exercise would reasonably use a lower value. Both inputs are editable, and $\bar{g}_P$ cancels out of the proportional gap in equation (16).

2.2 Frontier takeoff length, T

T measures the number of years from the start of the frontier automation wave to the modeled growth endpoint. The calendar date of AGI sits outside the model. In the falling-cost story, T is roughly the spread of task difficulty divided by the rate at which machine costs fall. It therefore bundles technical progress, differences across tasks, and the practical difficulty of economy-wide deployment.

[9] World Bank data make the heterogeneity plain. Over 2016–2024, real GDP per capita grew by roughly 1.9% a year in the United States, 5.3% in China, −0.3% in Sub-Saharan Africa, and −1.2% in Nigeria (World Bank, 2026).

The case for a short digital and research phase has become less speculative. METR's fits have the human-equivalent length of software tasks completed by frontier agents doubling roughly every four to seven months.[10] Jack Clark assigns more than 60% probability to no-human-involved AI R&D by the end of 2028 (Clark, 2026). I treat that as an informed personal forecast and do not feed it mechanically into T. Davidson et al. (2026) point the same way on feedbacks while staying deliberately silent on calendar timing.

So why 15 years? My rough story is that AI R&D and software-intensive sectors could move within a few years, while capital construction, power, factories, logistics, regulation, and organizational change stretch the transition. I allow roughly another decade for those slower margins. The phases overlap, so simply adding their lengths would overstate T. I hold 15 years loosely. The calculator exposes 5 to 30 years, and the slider goes wider.

2.3 Automation ceiling, $f_{\max}$

I treat $f_{\max}$ as a scenario parameter because I do not have a useful point estimate. The strongest takeoff story sets it to one. That could fail for technical reasons, including the blog post's explicit robotics assumption that machines eventually handle most physical production. It could also fail because the set of economically relevant tasks changes as people get richer.

Imas (2026) argues that if AI makes standardized commodities cheap, demand may move toward goods whose value is tied to a particular human being, such as care, teaching, live performance, or status. The calculator represents this possibility crudely through a lower ceiling on the task mass. It does not model shifting demand weights or a human-only sector whose relative price rises. The Imas essay therefore gives me a reason to vary $f_{\max}$, but no number with which to calibrate it.

I therefore interpret $f_{\max}$ as the fraction of the initially human task mass that is effectively automated by the end of the modeled transition. It can be below one because machines remain technically incapable of some tasks, but also because deployment is uneconomic, complementary capital is missing, regulation prevents automation, or consumers continue to value human production. The model treats this endpoint as fixed over the accounting window; it does not claim that it must remain the ultimate automation ceiling forever.

I keep $f_{\max} = 100\%$ central because the post is deliberately conditional on a strong takeoff scenario. Lower values show what happens when powerful AI capabilities do not translate into complete economy-wide automation during the detachment window. The calculator also shows 50%, 75%, 90%, and 100%, so the assumption is always visible.

2.4 Full-automation growth benchmark, g_{full}

The final input is the net growth rate in the limiting case in which the human-task bottleneck has disappeared. In the stripped-down accounting above,

$$g_{\text{full}} = sq_K - \delta.$$

This decomposition is useful because it shows what the parameter represents. But I do not attempt to calibrate s, q_K, and δ separately. Current investment and depreciation rates provide some evidence, but they need not remain fixed in an economy with radically higher returns to capital. More importantly, I do not have a credible estimate of q_K for an automated capital stock that includes robots, data centers, power systems, factories, mines, logistics networks, and the machinery needed to reproduce all of them.

I therefore pin down g_{full} directly rather than pretending to know each of its components. My central value comes from the numerical AGI scenarios in Korinek and Suh (2024). Once their economy enters an AK-like regime in

[10] The broad 2019–2025 fit doubles about every 196 days. The newer post-2023 estimate is faster, around 131 days. METR emphasizes that the estimate moves with the task suite, the tasks are concentrated in software engineering, machine learning, and cybersecurity, and measurements above 16 hours are currently unreliable (METR, 2026a,b).

which labor and capital are effectively substitutable, it converges to steady-state growth of approximately 18% per year.

Their calibration also makes clear that this number is not generated by machine productivity alone. In their notation, the productivity coefficient A plays the role of q_K. With an optimizing representative household, long-run growth is

$$g = \frac{A - \rho - \delta}{\eta},$$

where ρ is the rate of time preference and η is the curvature parameter in utility. Their choices $A = 0.5$, $\rho = 0.04$, $\delta = 0.10$, and $\eta = 2$ imply

$$g = \frac{0.50 - 0.04 - 0.10}{2} = 0.18.$$

The corresponding long-run investment share is 56%, which gives the same result through the accounting identity $g = s_\infty A - \delta$.

This illustrates why I treat g_{full} as a reduced-form parameter. The same 18% growth rate can be supported by different combinations of machine productivity, investment, and depreciation. For example, with a 25% investment share and 5% depreciation, it would require $q_K = 0.92$; Korinek and Suh instead combine $q_K = 0.5$ with a much higher endogenous investment share and 10% depreciation. The calculator does not try to distinguish between these decompositions.

I use 18% because it is a published full-automation benchmark that is closely matched to the scenario studied here—not because it is an empirical estimate or my precise forecast. Lower machine productivity, weaker investment incentives, or binding fixed factors could produce much lower growth. Conversely, sufficiently strong feedback from automated research could make a constant endpoint inappropriate because the growth rate itself would continue to rise (Davidson et al., 2026).

Equation (12) then combines this full-automation benchmark with the automation endpoint. With $g_0 = 2\%$, $g_{\text{full}} = 18\%$, and $\delta = 5\%$, $f_{\max} = 0.75$ implies $g_{\text{end}} = 12.1\%$, while $f_{\max} = 0.50$ implies $g_{\text{end}} = 7.7\%$. The full 18% benchmark is reached only when $f_{\max} = 1$.[11]

Lower-growth forecasts can reflect two different disagreements with the central scenario. Someone who thinks the economy will not approach complete effective automation during the window should lower $f_{\max}$ or increase T. Someone who accepts complete automation but thinks that machine capital will be expensive to reproduce, investment will remain limited, or fixed factors will bind should lower g_{full}. The model separates these two claims, although real-world forecasts may combine both.

2.5 The productive-adoption lag, D

Equation (8) interprets D as the time falling machine costs need to overcome a poorer economy's less favorable machine-to-labor cost ratio. I cannot estimate that wedge directly. The closest quantitative anchor is Comin and Mestieri (2018), who estimate diffusion paths for 25 major technologies across 139 countries.

Their estimates need two adjustments. Their headline lag starts the clock at invention, while this model starts it at frontier adoption. Also, D has to carry both the arrival of a technology and how intensively it is used once it arrives.

Start with the clock. Comin and Mestieri's adoption lag counts years between a technology's invention and its arrival in a country. Rich countries have positive lags too. Here $t = 0$ marks the start of the frontier takeoff, so the

[11] The calculator retains a 5% depreciation rate when interpolating gross growth between g_0 and g_{full}. Since g_{full} is pinned directly, this does not reproduce Korinek and Suh's complete parameterization. Changing δ changes the curvature of the interpolation but does not mechanically recalculate g_{full}. When $f_{\max} < 1$, a remaining human-task share also preserves diminishing returns in the bare task model, so g_{end} describes the modeled window rather than a fully specified long-run balanced-growth path.

relevant object is the difference between poor-country and rich-country lags. That is roughly the spread of their estimates.

Table 2: Arrival lags from invention for selected technologies

Technology	Mean	Median	P10–P90
Internet	6 years	7	3–9
Personal computers	14 years	14	11–17
Cellphones	14 years	16	11–19
All 25 technologies	42 years	37	8–82

Source: Comin and Mestieri (2018), Table 1.

Reading the tenth percentile as the fast adopters and the ninetieth as the slow ones, the frontier-relative spread is about 6 years for the internet and personal computers and about 8 years for cellphones. That spread is the useful arrival-margin anchor for D.[12]

Then there is intensity of use. Conditional on arrival, Comin and Mestieri measure how heavily a country uses the technology from the vertical shift of the same diffusion curve. Average intensity in non-Western countries is 47% of the Western level, and the 90th-to-10th percentile spread is roughly eightfold. Arrival lags narrowed over two centuries, while intensity gaps widened, especially for recent technologies. They attribute most of the income divergence since the mid-1800s to this second margin.

The D in this model shifts the whole productive-use path. It therefore has to absorb the arrival differential and the slower, shallower penetration that follows.[13]

The two adjustments pull in opposite directions and land somewhere sensible. I read the central value of 10 years as a 6–8 year arrival differential for digital-era technologies plus a few years for the intensity margin. The 15- and 20-year rows in the sensitivity table are live possibilities, and I put real weight on them.

Two cross-checks help. First, a tenfold cost-ratio wedge and a three-year halving time give $D = 3\log_2(10) \approx 10$ years through equation (9). Second, the contemporary bottleneck discussion, including scarce compute and Clark's economy-wide "Amdahl's law" examples (Clark, 2026), points to electricity, machinery, finance, complementary skills, and firm reorganization. None of those arrive at app-store speed.[14]

2.6 Evaluation horizon, H

H is the last year in the charts and welfare totals. Centrally, $H = T + D = 25$ years. The frontier boost reaches its endpoint at year T, and the lagged path receives the full boost at $T + D$. The central window therefore covers the whole period in which the lag is opening the gap. The blog post calls this the detachment window.

I stop there because the model has much less to say about what follows. Under the current specification, both developing-country paths eventually grow at $\bar{g}_P + B$. The proportional gap stops widening and does not mechanically close. Four continuations cover the main possibilities discussed in the blog post.

Future 1. The plateau. If growth stays at g_{end}, the gap persists at e^{BD}.

Future 2. The fade. If the boost is transient, equation (16) implies that the gap eventually closes. With

[12] Their lags are identified from the concave portion of historical diffusion curves and carry real estimation error. Their 139-country sample is also weighted differently from the calculator's LMIC-excluding-China block. Neither issue seems likely to change the order of magnitude.

[13] A permanent intensity shortfall behaves differently from a horizontal lag. In this model it would lower the developing-country ceiling, creating an $f_{\max}$ gap that shrinks the eventual boost. No value of D can reproduce that path exactly. Folding intensity into a few extra years of delay is the simple, conservative approximation.

[14] Early direct evidence points the same way even for the consumer margin. An IMF analysis of Windows-device telemetry finds that LLM use already lags substantially in poorer countries: https://www.imf.org/-/media/files/publications/dp/2026/english/upaiea.pdf.

$\int_0^\infty b(s)\,ds$ finite,

$$\lim_{t\to\infty} \int_{t-D}^{t} b(s)\,ds = 0.$$

The follower rides the tail of the boost after the frontier has slowed.

Future 3. **The singularity.** If the boost keeps growing, the gap keeps growing too.

Future 4. **The convergence.** If the follower outgrows the frontier after the window, the gap closes at high growth. This would need the post-window engine to run on non-rival ideas that can be copied freely. Owned capital comes with a price tag and a shipping delay.

A transition this large may also allocate more than income. It may determine ownership of self-replicating capital, military and political power, and the ability to write rules. Karnofsky's "competition" frame is one reason to be cautious about assuming a benign second phase (Karnofsky, 2021). The calculator prices phase one and stops. Welfare inside the window is already large, and the size of the gap at its end may shape whatever comes next.

3 From income paths to welfare and dollars

Let N_0 be the covered population, growing at rate n, so $N(t) = N_0 e^{nt}$. The welfare calculation uses log utility. Let V be the dollar value assigned to one unit of log income for one person-year, and let ρ be the welfare discount rate:

$$W(D) = V \int_0^H N(t) \ln\left(\frac{y^0(t)}{y^D(t)}\right) e^{-\rho t}\,dt. \tag{18}$$

The default $V = CG\$50{,}000$ makes a 1% income increase for one person for one year worth about CG\$500. This conversion is normative and editable. For intuition, the calculator restates W in lives-equivalent units by dividing by the value of a statistical life implied by the same conventions. The calculation uses CG\$100,000 per DALY and 33 DALYs per life, or CG\$3.3 million. The central welfare gap is equivalent to roughly 1.9 billion lives.

The calculator also reports the present value of income not realized under the lagged path. Let r be the dollar discount rate, set to 4% centrally while $\rho = 0$ for welfare:

$$S(D) = \int_0^H N(t) \left[y^0(t) - y^D(t)\right] e^{-rt}\,dt. \tag{19}$$

Both comparisons are between the two developing-country paths, so both equal zero when $D = 0$.[15]

[15] The browser code evaluates the integrals with 120 midpoint steps per year, updating income as $y(t+\Delta t) = y(t)\exp(g(t+\Delta t/2)\Delta t)$. Refining the grid moves the central estimate by far less than the displayed precision.

Table 3: Central inputs and outputs

Input or output	Central value	Interpretation
T	15 years	Length of the frontier automation transition
D	10 years	Effective delay in productive adoption
H	25 years	Central detachment window, $T + D$
g_0	2%	Frontier baseline growth before the takeoff
$f_{\max}$	100%	Share of remaining task mass effectively automated by the end of the transition
g_{full}	18%	Imported growth benchmark conditional on full automation
g_{end}	18%	Implied within-window endpoint under the central $f_{\max}$
δ	5%	Depreciation used in the growth interpolation
$\bar{g}_P$	2.5%	Separate developing-country baseline growth
N_0	5.4 billion	Rounded LMIC population excluding China
y_0	\$4,200	Rounded starting GDP per capita
V	CG\$50,000	Value of one log-income unit per person-year
ρ, r	0%, 4%	Welfare and dollar discount rates
$W(10)$	CG\$6.14 quadrillion	Welfare gap over 25 years
$W(10)$/CG\$3.3M	1.86 billion	Lives-equivalent restatement of the welfare gap
$S(10)$	\$1.25 quadrillion	Discounted income shortfall
$W(10) - W(9)$	CG\$381 trillion	Welfare value of one year faster adoption
$y^0(H)/y^{10}(H)$	4.95	Income ratio at the end of the window

The result moves a great deal with the main scenario choices. Table 4 changes one input at a time and holds the 25-year horizon fixed.

The last row reverses the order of the labels because a shorter takeoff creates a larger gap inside a fixed window. Section 1.2 also shows why the fast-takeoff, long-lag corner is internally strained. The same cost decline that compresses T tends to compress D. These sensitivities tell us much more than the second decimal place of the central estimate.

The population, income, and baseline-growth inputs are rounded from the World Bank's World Development Indicators. China is excluded because the motivating scenario treats it as part of the frontier (World Bank, 2026). These inputs set the scale of the calculation. I have not forecast them over the window.

Table 4: Selected one-at-a-time welfare sensitivities

Parameter varied	Lower case	Central case	Higher case
Adoption lag D	5 years: CG\$3.71 quadrillion	10 years: CG\$6.14 quadrillion	20 years: CG\$7.57 quadrillion
Automation ceiling $f_{\max}$	50%: CG\$2.28 quadrillion	75%: CG\$3.96 quadrillion	100%: CG\$6.14 quadrillion
Growth benchmark g_{full}	10%: CG\$3.17 quadrillion	18%: CG\$6.14 quadrillion	30%: CG\$10.42 quadrillion
Takeoff length T	30 years: CG\$2.43 quadrillion	15 years: CG\$6.14 quadrillion	5 years: CG\$9.00 quadrillion

4 Limitations and scope

1. **I do not estimate the effective automation endpoint.** $f_{\max}$ makes the central full-automation assumption visible. Current evidence does not tell us whether 60%, 90%, or 100% of the initially human task mass will be effectively automated during the modeled window.
2. **A fixed task share misses structural change.** If automated goods become cheap and demand shifts toward relational or human-intrinsic services, the expenditure weight on unautomated tasks may rise. The $f_{\max}$

sensitivity is only a rough proxy for the mechanism in Imas (2026).

3. **I import rather than estimate the full-automation growth benchmark.** The central $g_{\text{full}} = 18\%$ comes from the numerical AGI calibration in Korinek and Suh (2024). It is not a general implication of full automation. It bundles assumptions about the productivity of reproducible capital, depreciation, household preferences, and endogenous investment. Fixed-factor bottlenecks could make the rate much lower; strong research feedbacks could make a constant endpoint inappropriate because growth would continue accelerating.

4. **Falling costs motivate the takeoff curve.** A logistic path follows from a log-logistic spread of task difficulty and roughly exponential machine-cost declines. The calculator leaves the difficulty distribution, λ, research productivity, investment behavior, fixed factors, and the joint evolution of capital and labor productivity unestimated.

5. **Partial automation makes the long run especially uncertain.** If a human-task bottleneck remains, the task model keeps diminishing returns to capital. The endpoint describes the modeled window. Elevated growth may or may not persist after it.

6. **One lag bundles several margins.** Arrival, profitability, capital installation, organizational adoption, and intensity of use are different things. Section 1.2 shows how wages and deployment costs generate a horizontal lag. Section 2.5 folds the intensity margin into extra years. A permanent intensity shortfall would instead lower the ceiling, and D alone cannot express it. The derivation also assumes a common task mix and a common machine-cost decline. When those differ, the effective lag drifts over the transition.

7. **Historical lags come from technologies with ordinary payoffs.** Comin and Mestieri's lags are equilibrium outcomes under the adoption incentives of tractors, telephones, and the internet. A growth explosion could change those incentives sharply. Huge returns should pull adoption forward when incentives are the binding constraint. They may do little when the constraints are capacity, scarce dollars, power, skills, or a world supply of machines price-rationed to high-wage markets first. The threshold model offers a partial defense because adoption timing depends on the cost crossing. That feature also reflects the simplifying assumption that fixed adoption costs are absent.

8. **The model assumes that the same AI boost eventually arrives.** Both counterfactual paths receive the same proportional boost. Persistent differences in complementary capital, institutions, ownership, or sectoral composition could make the eventual boost smaller as well as later. The $f_{\max}$ and D sensitivities capture parts of that concern, but not all of it.

9. **General-equilibrium effects are absent.** The model omits trade, capital flows, migration, exchange rates, ownership of AI capital, political responses, and terms-of-trade effects. Any of these could narrow or widen the gap.

10. **The representative developing-country block hides distribution.** GDP per capita can rise while gains remain very uneven across countries and people.

11. **The model stops at H.** Once both paths reach the endpoint, the specification makes the gap stop widening. It leaves eventual convergence and permanent lock-in open.

The central number is too fragile for a literal bet. The scale arithmetic is harder to escape. Fast frontier growth turns even a modest delay into a large proportional income gap. Partial automation shrinks the gap, though it can still remain large. Multiply that gap by billions of people and several years and the welfare stakes become enormous very quickly.

A Derivations

A.1 Capital growth and the interpolation in equation (11)

Define capital per effective worker $x \equiv K/(AL)$ and the entropy term $E(f) \equiv -f \ln f - (1-f)\ln(1-f)$. Equation (4) can then be written $Y = ALe^{E(f)}x^f$. With $\dot{K} = sY - \delta K$,

$$g_K + \delta = s\frac{Y}{K} = se^{E(f)}x^{-(1-f)}. \tag{20}$$

As f rises, the drag exponent $1-f$ shrinks. Reproducible capital is less constrained by the tasks that still require people. Taking logs gives

$$\ln(g_K + \delta) = \ln s + E(f) - (1-f)\ln x. \tag{21}$$

Let $u \equiv (f - f_0)/(1 - f_0)$ be the fraction of the task mass that was human at the start of the wave and has since been automated. This is only a rescaling of f, not a second automation variable. Both x and the entropy term $E(f)$ generally change during the transition, so the task model does not by itself imply a log-linear relationship between automation and gross growth. I therefore use an endpoint-pinned approximation: I approximate $\ln(g_F + \delta)$ by the straight line in u connecting the observed baseline growth rate and the imported full-automation benchmark. This discards the curvature generated by $E(f)$ and changes in x, but provides a transparent mapping with no additional free curvature parameter. Pinning the approximation to g_0 at $u = 0$ and g_{full} at $u = 1$ gives

$$\ln(g_F + \delta) = (1-u)\ln(g_0 + \delta) + u\ln(g_{\text{full}} + \delta). \tag{22}$$

Substituting $u = (f_F - f_0)/(1 - f_0)$ gives equation (11), and substituting $u = f_{\max}$ gives equation (12). The pinning removes the need to estimate f_0 because observed g_0 already contains whatever automation exists today. The implied curvature over the full transition is $\beta = \ln[(g_{\text{full}} + \delta)/(g_0 + \delta)]$, so the endpoints determine the mapping.

Capital growth and GDP growth differ during the transition. Exact differentiation of the production function gives

$$g_Y = fg_K + (1-f)g_{AL} + \dot{f}\ln\left(\frac{K/f}{AL/(1-f)}\right), \tag{23}$$

where g_{AL} is growth of effective labor and the last term is the direct effect of moving the task boundary. A full calibration would need much more structure. The endpoint-pinned curve is the transparent shortcut I use here.

The normalization in equation (4) suppresses q_K. Restoring it would scale the gross return to capital, and at full automation the accounting identity becomes $g_K + \delta = sq_K$. Since both the investment share and q_K may change sharply in the scenario of interest, I do not infer g_{full} from current values of either one. Instead, I import a net endpoint growth rate from Korinek and Suh's numerical full-automation scenario and use equation (11) only to interpolate between the observed baseline and that endpoint.

A.2 Why a logistic automation path

Take logs of the threshold condition in equation (6). Task i is automated when $\ln \kappa_i \le \lambda t + \ln c_j^L - \ln \omega_j$. If log difficulty is logistically distributed with midpoint μ and scale ν, the machine-task share is

$$f_j(t) = \sigma\left(\frac{\lambda t + \ln c_j^L - \ln \omega_j - \mu}{\nu}\right), \tag{24}$$

an S-curve in time. The time to move from 5% to 95% of the relevant task mass is

$$T_{5\text{–}95} = \frac{2\nu \ln 19}{\lambda} \simeq \frac{5.9\nu}{\lambda}. \tag{25}$$

A short takeoff can therefore come from fast machine-cost declines, a narrow spread of task difficulty, or both. Equation (7) truncates and rescales this logistic over the modeled transition, starts it at f_0, and caps it at $f_\infty = f_0 + (1 - f_0) f_{\max}$. The factor 6 is a rounded version of $2 \ln 19 \approx 5.9$.

Set the two countries' logistic arguments equal and solve for the time shift. This gives equation (8). The difficulty distribution cancels, which is why D does not depend on it. Dividing equation (8) by equation (25) gives

$$\frac{D}{T_{5\text{–}95}} = \frac{\ln[(\omega_P / c_P^L)/(\omega_F / c_F^L)]}{2\nu \ln 19}.$$

The ratio of the lag to the takeoff length is independent of λ.

References

Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*, Vol. 4B, edited by Orley Ashenfelter and David Card, 1043–1171. Amsterdam: Elsevier. doi:10.1016/S0169-7218(11)02410-5.

Acemoglu, Daron, and Pascual Restrepo. 2018. "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review* 108(6): 1488–1542. doi:10.1257/aer.20160696.

Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones. 2019. "Artificial Intelligence and Economic Growth." In *The Economics of Artificial Intelligence: An Agenda*, edited by Ajay Agrawal, Joshua Gans, and Avi Goldfarb. Chicago: University of Chicago Press. Working-paper version: doi:10.3386/w23928.

Clark, Jack. 2026. "Import AI 455: Automating AI Research." *Import AI*, May 4. `https://jack-clark.net/2026/05/04/import-ai-455-automating-ai-research/`.

Comin, Diego, and Marti Mestieri. 2018. "If Technology Has Arrived Everywhere, Why Has Income Diverged?" *American Economic Journal: Macroeconomics* 10(3): 137–178. doi:10.1257/mac.20150175.

Davidson, Tom, Basil Halperin, Thomas Houlden, and Anton Korinek. 2026. "When Does Automating AI Research Produce Explosive Growth? Feedback Loops in Innovation Networks." NBER Working Paper 35155. doi:10.3386/w35155.

Erdil, Ege, Andrei Potlogea, Tamay Besiroglu, Edu Roldan, Anson Ho, Jaime Sevilla, Matthew Barnett, Matej Vrzla, and Robert Sandler. 2025. "GATE: An Integrated Assessment Model for AI Automation." arXiv:2503.04941.

Imas, Alex. 2026. "What Will Be Scarce? The Economics of Structural Change and the Post-Commodity Future of Work." *Ghosts of Electricity*, April 14. `https://aleximas.substack.com/p/what-will-be-scarce`. Companion note: "Structural Change, AI, and the Relational Sector." `https://www.aleximas.com/s/Technical-Note-Structural-Change.pdf`.

Karnofsky, Holden. 2021. "How to Make the Best of the Most Important Century?" *Cold Takes*, September 14. `https://www.cold-takes.com/making-the-best-of-the-most-important-century/`.

Korinek, Anton, and Donghyun Suh. 2024. "Scenarios for the Transition to AGI." NBER Working Paper 32255. doi:10.3386/w32255.

METR. 2026a. "Task-Completion Time Horizons of Frontier AI Models." Updated May 8, 2026. `https://metr.org/time-horizons/`.

METR. 2026b. "Time Horizon 1.1." January 29, 2026. `https://metr.org/blog/2026-1-29-time-horizon-1-1/`.

Trammell, Philip, and Anton Korinek. 2026. "Economic Growth under Transformative AI." NBER Working Paper 31815, revised April 2026. doi:10.3386/w31815.

World Bank. 2026. *World Development Indicators*. Population: `https://data.worldbank.org/indicator/SP.POP.TOTL`. GDP: `https://data.worldbank.org/indicator/NY.GDP.MKTP.CD`. GDP-per-capita growth: `https:`

//data.worldbank.org/indicator/NY.GDP.PCAP.KD.ZG. Population growth: https://data.worldbank.org/indicator/SP.POP.GROW. Accessed July 2026.

Zeira, Joseph. 1998. "Workers, Machines, and Economic Growth." *Quarterly Journal of Economics* 113(4): 1091–1117. doi:10.1162/003355398555847.