"We Won't be Missed: Work and Growth in the AGI World" Pascual Restrepo. Forthcoming in The Economics of Transformative AI, edited by A. Agrawal, E. Brynjolfsson, and A. Korinek. University of Chicago Press.

Online Appendix: Proofs and Derivations (Not for Publication)

My derivations use two technical lemmas from convex analysis.

Lemma 1. *For any set of computational requirements* $\{\alpha(\omega)\}$ *(including* $\alpha(\omega) = \infty$*), the function*

$$M(Q, \{H(s)\}_{s\in\mathcal{S}}) = \max_{\{Q(\omega),L(\omega)\}} F(\{L(\omega) + \frac{1}{\alpha(\omega)} Q(\omega)\}) \text{ st: } \sum_{\omega} Q(\omega) \leq Q, \quad \sum_{\omega\in\Omega(s)} L(\omega) \leq H(s) \text{ for all } s \in \mathcal{S} \tag{5}$$

is jointly concave in $Q, \{H(s)\}$.

Proof. Consider $y_1 = (Q_1, \{H_1(s)\}_{s\in\mathcal{S}})$ and $y_2 = (Q_2, \{H_2(s)\}_{s\in\mathcal{S}})$, and let $x_1 = (\{L_1(\omega), Q_1(\omega)\}_{\omega\in\Omega})$ and $x_2 = (\{L_2(\omega), Q_2(\omega)\}_{\omega\in\Omega})$ be corresponding maximizers. For any $\lambda \in [0, 1]$, $\lambda x_1 + (1 - \lambda)x_2 \in \mathbb{X}$ is a feasible allocation for an endowment level of $\lambda y_1 + (1 - \lambda)y_2$, since all resource constraints are affine. This implies:

$$\begin{aligned} M(\lambda y_1 + (1-\lambda)y_2) &\geq F(\lambda x_1 + (1-\lambda)x_2) \\ &\geq \lambda F(x_1) + (1-\lambda)F(x_2) \\ &= \lambda M(x_1) + (1-\lambda)M(x_2) \end{aligned}$$

and M is concave. □

Lemma 2. *Let* $A(Q, \{H(s)\}_{s\in\mathcal{S}}) \geq 0$ *be the Lagrange multiplier on computing resources in* (5). *The multiplier is finite and (weakly) decreases in* Q*, while the function*

$$M(Q, \{H(s)\}_{s\in\mathcal{S}}) - A(Q, \{H(s)\}_{s\in\mathcal{S}})\, Q$$

(weakly) increases in Q *for all* $Q \geq 0$.

Proof. The envelope theorem says that $A(Q, \{H(s)\}_{s\in\mathcal{S}})$ is a super-gradient of M with respect to Q. That is, the line going through $M(Q, \{H(s)\}_{s\in\mathcal{S}})$ with slope $A(Q, \{H(s)\}_{s\in\mathcal{S}})$ is above M. The concavity of M then implies that $A(Q, \{H(s)\}_{s\in\mathcal{S}})$ weakly decreases in Q.

Therefore, for any $Q_1 > Q_0$, we have

$$\frac{M(Q_1, \{H(s)\}_{s\in\mathcal{S}}) - M(Q_0, \{H(s)\}_{s\in\mathcal{S}})}{Q_1 - Q_0} \;\geq\; A(Q_1, \{H(s)\}_{s\in\mathcal{S}}).$$

Rearranging terms

$$M(Q_1, \{H(s)\}_{s\in\mathcal{S}}) - A(Q_1, \{H(s)\}_{s\in\mathcal{S}})\, Q_1 \;\geq\; M(Q_0, \{H(s)\}_{s\in\mathcal{S}}) - A(Q_1, \{H(s)\}_{s\in\mathcal{S}})\, Q_0.$$

From the concavity of M we conclude that $A(Q_0, \{H(s)\}_{s\in\mathcal{S}}) \geq A(Q_1, \{H(s)\}_{s\in\mathcal{S}})$, which implies that we can substitute $A(Q_0, \{H(s)\}_{s\in\mathcal{S}})$ for $A(Q_1, \{H(s)\}_{s\in\mathcal{S}})$ on the right and preserve the inequality:

$$M(Q_1, \{H(s)\}_{s\in\mathcal{S}}) - A(Q_1, \{H(s)\}_{s\in\mathcal{S}})\, Q_1 \;\geq\; M(Q_0, \{H(s)\}_{s\in\mathcal{S}}) - A(Q_0, \{H(s)\}_{s\in\mathcal{S}})\, Q_0.$$

This shows that $M(Q, \{H(s)\}_{s\in\mathcal{S}}) - A(Q, \{H(s)\}_{s\in\mathcal{S}})\, Q$ is increasing in Q for $Q \geq 0$ as wanted. □

Proofs of the propositions in the main text

In the proofs, I assume the economy is competitive and maximizes output at each point in time subject to its computational resources Q_t and skill endowment $H(s)$. The derivations assume that this problem has a unique solution (in terms of $X_t(\omega)$) and that the resulting output is a differentiable function of Q_t and human skills $H(s)$, with continuous derivatives.

Proof of Proposition 1. We first show that output Y_t must grow in an unbounded way. Because the economy is organized efficiently,

$$Y_t \;\geq\; F(\{L_t(\omega) + \frac{1}{\alpha_t(\omega)}\,\frac{Q_t}{N}\}),$$

where N is the size of the set Ω. Because output increases in $X_t(\omega)$, this also implies

$$Y_t \;\geq\; F(\{\frac{1}{\alpha_t(\omega)}\,\frac{Q_t}{N}\}) \;=\; F(\{\frac{1}{\alpha_t(\omega)}\,\frac{1}{N}\})\, Q_t \;\geq\; \underbrace{F(\{\frac{1}{\alpha(\omega)}\,\frac{1}{N}\})}_{=\underline{A}}\, Q_t.$$

The first equality uses the fact that F is constant returns to scale. The second inequality uses the fact that $\alpha_t(\omega)$ decreases and converges to a finite constant. This shows that $Y_t \geq \underline{A}\ Q_t$, for some $\underline{A} > 0$, which shows that output grows in an unbounded way.

Let $\{X_t(\omega)\}$ denote an optimal production plan. The above observation shows that $F(\{X_t(\omega)\})$ is unbounded.

Suppose ω_0 is a bottleneck. The definition of bottlenecks implies that either i. $X_t(\omega_0)$ is unbounded or ii. $F_{\omega_0}(\{X_t(\omega)\})$ is unbounded along an optimal production plan.

In case i, we conclude that ω_0 must be produced with AI systems, since otherwise $X_t(\omega_0) = L_t(\omega_0)$ would be bounded.

In case ii, we conclude that ω_0 must also be produced with AI systems. Denote by A_t the marginal product of compute Q_t at time t (that is, the Lagrange multiplier of Q at time t). Because of CRS and marginal diminishing returns to work, A_t eventually decreases in time. Specifically, once $\alpha_t(\omega) = \alpha(\omega)$ for all ω, the only source of technological progress are increases in Q_t, which weakly reduce A_t (from Lemma 2). As a result, we reach a time when

$$F_{\omega_0}(\{X_t(\omega)\}) > A_t\ \alpha_t(\omega_0),$$

and reallocating one unit of compute from other uses to producing $X_t(\omega_0)$ raises output, implying that ω_0 must be eventually produced with AI systems. □

Proof of Proposition 2. Fix an allocation of labor $\{L(\omega)\}_{\omega\in\Omega}$. I first characterize the limit behavior of the production frontier of the economy for this allocation, as compute Q increases and $\alpha_t(\omega) \to \alpha(\omega)$. I then use this to characterize the behavior of the production frontier in terms of skills, once we allow labor to be reallocated within $\Omega(s)$.

For large values of Q, the maximum output we can attain at this allocation is

$$\begin{aligned} M(\widetilde{Q}, \{L(\omega)\}_{\omega\in\Omega_{\mathcal{N}}}) \ = \ & \max_{\{\widetilde{Q}(\omega)\}_{\omega\in\Omega_{\mathcal{A}}}} \ F\big(\{\frac{1}{\alpha(\omega)}\ \widetilde{Q}(\omega)\}_{\omega\in\Omega_{\mathcal{A}}}, \{L(\omega)\}_{\omega\in\Omega_{\mathcal{N}}}\big) \\ & \text{s.t: } \sum_{\omega\in\Omega_{\mathcal{A}}} \widetilde{Q}(\omega) \leq \underbrace{Q + \sum_{\omega\in\Omega_{\mathcal{A}}} \alpha(\omega)\ L(\omega)}_{=\widetilde{Q}}. \end{aligned}$$

Here $\Omega_{\mathcal{A}}$ is the set of work produced by AI systems and $\Omega_{\mathcal{N}}$ the set produced exclusively by workers in the long run (given the fixed allocation of labor and large quantities of compute Q). The function

M is concave (this is a special case of Lemma 1) and features constant-returns to scale.

I first characterize the long-run behavior of M as compute $\widetilde{Q}$ increases.

Let $A(\widetilde{Q}, \{L(\omega)\}_{\omega \in \mathcal{N}})$ denote the Lagrange multiplier of compute. Since M is assumed differentiable, the envelope theorem implies its derivative with respect to Q satisfies

$$M_Q(\widetilde{Q}, \{L(\omega)\}_{\omega \in \Omega_{\mathcal{N}}}) \;=\; A(\widetilde{Q}, \{L(\omega)\}_{\omega \in \Omega_{\mathcal{N}}}).$$

Define

$$V(\widetilde{Q}, \{L(\omega)\}_{\omega \Omega_{\in} \mathcal{N}}) \;=\; M(\widetilde{Q}, \{L(\omega)\}_{\omega \in \Omega_{\mathcal{N}}}) - A(\widetilde{Q}, \{L(\omega)\}_{\omega \in \Omega_{\mathcal{N}}})\, \widetilde{Q}.$$

Lemma 2 implies that V increases in $\widetilde{Q}$.

In addition, the fact that work in Ω_N is not automated (even though it could be) implies that V is bounded from above. To show this, note that constant returns in M plus Euler's theorem imply

$$V(\widetilde{Q}, \{L(\omega)\}_{\omega \in \Omega_{\mathcal{N}}}) = \sum_{\omega_0 \in \Omega_{\mathcal{N}}} M_{\omega_0}(\widetilde{Q}, \{L(\omega)\}_{\omega \in \mathcal{N}})\, L(\omega).$$

Since $\omega_0 \in \Omega_{\mathcal{N}}$ is not automated, we must have that for large $\widetilde{Q}$,

$$M_{\omega_0}(\widetilde{Q}, \{L(\omega)\}_{\omega \in \Omega_{\mathcal{N}}}) < A(\widetilde{Q}, \{L(\omega)\}_{\omega \in \Omega_{\mathcal{N}}})\, \alpha(\omega_0).$$

Else, we would get more output by reallocating some compute to producing ω_0 (the marginal benefit would exceed the opportunity cost of compute). This implies that

$$V(\widetilde{Q}, \{L(\omega)\}_{\omega \in \Omega_{\mathcal{N}}}) \;<\; A(\widetilde{Q}, \{L(\omega)\}_{\omega \in \Omega_{\mathcal{N}}}) \sum_{\omega \in \Omega_{\mathcal{N}}} \alpha(\omega)\, L(\omega),$$

and V is bounded from above. We conclude that, as compute grows in time,

$$V(\widetilde{Q}, \{L(\omega)\}_{\omega \in \Omega_{\mathcal{N}}}) \to V(\{L(\omega)\}_{\omega \in \Omega_{\mathcal{N}}})$$

for some finite limit function V that only depends on $\{L(\omega)\}_{\omega \in \Omega_{\mathcal{N}}}$.

I now show that

$$\lim_{\widetilde{Q} \to \infty} A(\widetilde{Q}, \{L(\omega)\}_{\omega \in \Omega_{\mathcal{N}}}) \to A,$$

for some constant $A > 0$. Lemma 1 shows that $A(\widetilde{Q}, \{L(\omega)\}_{\omega\in\Omega_{\mathcal{N}}})$ is non-negative and decreasing in $\widetilde{Q}$ (since M is concave). From

$$\frac{M(\widetilde{Q}, \{L(\omega)\}_{\omega\in\Omega_{\mathcal{N}}})}{\widetilde{Q}} = A(\widetilde{Q}, \{L(\omega)\}_{\omega\in\Omega_{\mathcal{N}}}) + \frac{V(\widetilde{Q}, \{L(\omega)\}_{\omega\in\Omega_{\mathcal{N}}})}{\widetilde{Q}},$$

we conclude that the limit of $A(\widetilde{Q}, \{L(\omega)\}_{\omega\in\Omega_{\mathcal{N}}})$ as $\widetilde{Q}$ increases must be positive. Otherwise, the right side of this equality would be zero (recall that V is bounded) while the left would be positive (recall that output is bounded from below by $\underline{A}\ \widetilde{Q}$, as shown in the proof of Proposition 1).

The above derivations show that, in the long run, and as compute increases

$$M(\widetilde{Q}, \{L(\omega)\}_{\omega\in\Omega_{\mathcal{N}}}) \rightarrow A\ \widetilde{Q} + V(\{L(\omega)\}_{\omega\in\Omega_{\mathcal{N}}}),$$

for some positive constant $A > 0$ and some function V, bounded from above by

$$V(\{L(\omega)\}_{\omega\in\Omega_{\mathcal{N}}}) < \sum_{\omega\in\Omega_{\mathcal{N}}} A\ \alpha(\omega)\ L(\omega). \tag{6}$$

I now turn to the long-run production frontier of the economy in terms of compute and skills. The maximum output we can get from a stock of compute Q_t and skills $\{H(s)\}_{s\in\mathcal{S}}$ is

$$M^*(Q, \{H(s)\}_{s\in\mathcal{S}}) = \max_{\{L(\omega)\}_{\omega\in\Omega}} M(\widetilde{Q}, \{L(\omega)\}_{\omega\in\Omega_{\mathcal{N}}}) \text{ st: } \widetilde{Q} = Q + \sum_{\omega} \alpha(\omega)\ L(\omega),$$
$$\sum_{\omega\in\Omega(s)} L(\omega) \leq H(s) \text{ for all } s \in \mathcal{S}$$

Using the limit behavior of M, we can rewrite this as

$$M^*(Q, \{H(s)\}_{s\in\mathcal{S}}) = \max_{\{L(\omega)\}_{\omega\in\Omega}} A\left(Q_t + \sum_{\omega\in\Omega_{\mathcal{A}}} \alpha(\omega)\ L_t(\omega)\right) + V(\{L(\omega)\}_{\omega\in\mathcal{N}})$$
$$\text{st: } \sum_{\omega\in\Omega(s)} L(\omega) \leq H(s) \text{ for all } s \in \mathcal{S}. \tag{7}$$

One can distinguish between two cases for each skill. For group $\mathcal{A}$, at least one type of work in $\arg\max_{\omega\in\Omega(s)} \alpha(\omega)$ lies in $\Omega_{\mathcal{A}}$ (i.e., the "best" use for s is in work that is also produced by AI). For these skills, the right-side of (7) is maximized by allocating all labor of skill s to work in

$\arg\max_{\omega \in \Omega(s)} \cap \Omega_{\mathcal{A}}$, resulting in an output of $A\ \mathrm{CEU}(s)\ H(s)$. Note that these skills are not used for any work in $\Omega_{\mathcal{N}}$, as this would yield a lower marginal product of at most $A\ \alpha(\omega_0)$ (from the fact that no compute is allocated to this work).

For group $\mathcal{N}$, all work in $\arg\max_{\omega \in \Omega(s)} \alpha(\omega)$ lies in $\Omega_{\mathcal{N}}$. These skills account for all labor used in $\Omega_{\mathcal{N}}$ and some of the labor used in automated work. Therefore, we can write the contribution of these skills to output as

$$N(\{H(s)\}_{s \in \mathcal{N}}) = \max_{\{L(\omega)\}_{\omega \in \Omega(s): s \in \mathcal{N}}} \sum_{s \in \mathcal{N}} \sum_{\omega \in \Omega(s) \cap \Omega_{\mathcal{A}}} A\ \alpha(\omega)\ L(\omega) + V(\{L(\omega)\}_{\omega \in \mathcal{N}})$$
$$\text{st: } \sum_{\omega \in \Omega(s)} L(\omega) \leq H(s) \text{ for all } s \in \mathcal{N}. \tag{8}$$

This problem admits three classes of solutions: skills $s \in \mathcal{N}$ produce supplementary non-automated work, automated work outside $\Omega^*(s)$, or both.

These derivations show that we can write

$$M^*(Q, \{H(s)\}_{s \in \mathcal{S}}) = A \left(Q_t + \sum_{s \in \mathcal{A}} \mathrm{CEU}(s)\ H(s) \right) + N(\{H(s)\}_{s \in \mathcal{N}}), \tag{9}$$

with N defined in (8). Moreover, the bound on V in (6) implies that

$$N(\{H(s)\}_{s \in \mathcal{N}}) < \sum_{s \in \mathcal{N}} \left(\sum_{\omega \in \Omega(s) \cap \Omega_{\mathcal{A}}} A\ \alpha(\omega)\ L(\omega) + \sum_{\omega \in \Omega(s) \cap \Omega_{\mathcal{N}}} A\ \alpha(\omega)\ L(\omega) \right) \leq \sum_{s \in \mathcal{N}} A\ \mathrm{CEU}(s)\ H(s)$$

as wanted. □

Proof of Proposition 3. Follows from the expression for output in Proposition 2. □

Proof of Proposition 4. The expression for output in Proposition 2 shows that the marginal product of skill $s \in \mathcal{A}$ (and so its wage) is

$$W(s) = A\ \mathrm{CEU}(s).$$

For all other skills, wages converge to

$$W(s) = \max \left\{ \{A\ \alpha(\omega_0)\}_{\omega_0 \in \Omega(s) \cap \Omega_{\mathcal{A}}}, \{M_{\omega_0}(\widetilde{Q}, \{L(\omega)\}_{\omega \in \Omega_{\mathcal{N}}})\}_{\omega_0 \in \Omega(s) \cap \Omega_{\mathcal{N}}} \right\} < A\ \mathrm{CEU}(s).$$

This uses the fact that, for $s \in \mathcal{N}$, $\alpha(\omega_0) < \text{CEU}(s)$ for all $\omega_0 \in \Omega(s) \cap \Omega_{\mathcal{A}}$ (by definition of the set $\mathcal{N}$), and $M_{\omega_0}(M_{\omega_0}(\widetilde{Q}, \{L(\omega)\}_{\omega \in \Omega_{\mathcal{N}}})) < A\ \alpha(\omega_0) \leq A\ \text{CEU}(s)$ (as otherwise, ω_0 should be produced with compute). □

Proof of Proposition 5. The labor share is

$$\frac{\sum_{s \in \mathcal{A}} \text{CEU}(s)\ H(s)\ +\ \sum_{s \in \mathcal{N}} \widetilde{\text{CEU}}(s)\ H(s)}{Q_t + \sum_{s \in \mathcal{A}} \text{CEU}(s)\ H(s)\ +\ \sum_{s \in \mathcal{N}} \widetilde{\text{CEU}}(s)\ H(s)}.$$

This is bounded from above by

$$\frac{\sum_s \text{CEU}(s)\ H(s)}{Q_t + \sum_s \text{CEU}(s)\ H(s)},$$

which converges to zero as Q_t expands over time. □

Proof of Proposition 6. The argument is given in the text. □

As explained in the text, the results for the transition assume there is $q-$complementarities across all forms of work. Formally, what I mean by this is:

Assumption 1 ($q-$complements assumption)**.** *Let $\mathcal{P}_1, \ldots, \mathcal{P}_N$ be an arbitrary partition of Ω, such that all $\omega \in \mathcal{P}_i$ are allocated efficiently across an input in total supply X_i, with productivity $\beta(\omega) > 0$ in each $\omega \in \mathcal{P}_i$. Then, the marginal value of X_i is increasing and continuous in all X_j for $j \neq i$, and decreasing and continuous in X_i.*

Proof of Proposition 7. Define

$$M(Q, \{H(s)\}_{s \in \mathcal{S}}; t)\ =\ \max F\Big(\{L(\omega) + \frac{1}{\alpha_t(\omega)} Q(\omega)\}\Big)$$
$$\text{st: } \sum_{\omega} Q(\omega) \leq Q \ \text{ and } \ \sum_{s \in \Omega(s)} L(\omega) \leq H(s).$$

Lemma 1 implies that $M(Q, \{H(s)\}_{s \in \mathcal{S}}; t)$ is concave. Moreover, it is assumed in the following derivations that M is differentiable and has continuous derivatives.

Output at time t can be written as

$$Y_t\ =\ M(Q_t, \{H(s)\}_{s \in \mathcal{S}}; t),$$

while the MRT of compute into output is

$$A_t \;=\; M_Q(Q_t, \{H(s)\}_{s\in\mathcal{S}}; t).$$

Finally, from constant-returns to scale for M and Eulers' theorem, we can write the wage bill as

$$W_t \;=\; M(Q_t, \{H(s)\}_{s\in\mathcal{S}}; t) - M_Q(Q_t, \{H(s)\}_{s\in\mathcal{S}}; t)\; Q_t.$$

Let $T(\omega_1) < \ldots < T(\omega_N)$ be the times at which work becomes automatable. The functions Y_t, A_t, W_t are continuous in $(0, T(\omega_1))$, $(T(\omega_1), T(\omega_2))$, ..., $(T(\omega_N), \infty)$, since during these intervals, only Q_t changes, this change is continuous (recall that Q_t is assumed continuous in time), and $M(Q, \{H(s)\}_{s\in\mathcal{S}}; t)$ and $M_Q(Q, \{H(s)\}_{s\in\mathcal{S}}; t)$ are continuous in Q for all t.

Moreover, the assumption that compute binds implies that Y_t, A_t, W_t are continuous at $T(\omega)$ for all ω, since the technology is not implemented at this time. This shows that the aggregates Y_t, A_t, W_t have no jumps or discontinuities.

Next, I show that A_t is decreasing and W_t is increasing in time. The fact that A_t is decreasing follows from the fact that A_t does not jump at $T(\omega)$ and is otherwise decreasing in Q_t (from Lemma 2). Likewise, W_t does not jump at $T(\omega)$ and is otherwise increasing in Q_t (from Lemma 2).

I now turn to wages and the allocation of labor. Wages are given by

$$W_t(s) \;=\; M_{H(s)}(Q_t, \{H(s)\}_{s\in\mathcal{S}}; t).$$

As a function of time, $W_t(s)$ is continuous in $(0, T(\omega_1))$, $(T(\omega_1), T(\omega_2))$, ..., $(T(\omega_N), \infty)$, since during these intervals, only Q_t changes, this change is continuous (recall that Q_t is assumed continuous in time), and $M_{H(s)}(Q, \{H(s)\}_{s\in\mathcal{S}}; t)$ is continuous in Q. Moreover, the assumption that compute binds implies that $M_{H(s)}(Q, \{H(s)\}_{s\in\mathcal{S}}; t)$ is continuous at $T(\omega)$ for all ω, since algorithmic advances have no consequences when they arrive.

Let's now characterize the dynamics of employment and wages for skill s. Let's order the work in $\Omega(s)$ according to the time it becomes automated $\tilde{T}(\omega)$. Assume there are $N(s)$ elements in $\Omega(s)$, automated at dates

$$\tilde{T}(\omega_1) \le \tilde{T}(\omega_2) \le \ldots, \tilde{T}(\omega_{N(s)}),$$

and suppose that $N(s) \geq 2$ (the case with $N(s) = 1$ is considered below).

Let's analyze first what happens when ω_1 is automated. This takes place at a time $\tilde{T}(\omega_1) > T(\omega_1)$, such that

$$W_t(s) \;=\; A_t \;\alpha(\omega_1).$$

I now show that for some interval of time after $\tilde{T}(\omega_1)$, $L_t(\omega_1)$ (i) reaches zero after some finite period of transition, (ii) decreases during this period, and (iii) does so continually.

For (i), suppose by way of contradiction that $L_t(\omega_1) > 0$ for all $t > \tilde{T}(\omega_1)$. Because labor is being used for ω_1, we must have $W_t(s) = A_t \;\alpha(\omega_1)$ for all $t > \tilde{T}(\omega_1)$, since wages must be equal to the cost of producing ω_1 with AI systems. Let's now consider ω_2. We have two potential cases:

1. If $\alpha(\omega_2) < \alpha(\omega_1)$, $T(\omega_2)$ must be greater than $\tilde{T}(\omega_1)$ or ω_2 would have been automated first. However, note that this implies that at time $T(\omega_2)$, we have $W_t(s) = A_t \;\alpha(\omega_1) > A_t \;\alpha(\omega_2)$, which implies AI systems would be adopted immediately, contradicting the premise that compute binds at all points in the transition.

2. If $\alpha(\omega_2) > \alpha(\omega_1)$, then $W_t(s) = A_t \;\alpha(\omega_1) < A_t \;\alpha(\omega_2)$ for all $t > \tilde{T}(\omega_1)$, but this would imply ω_2 is never automated, contradicting the assumption that all work is a bottleneck.

This shows that, if there is some other work $\omega_2 \in \Omega(s)$ that has not been automated at time $\tilde{T}(\omega_1)$, then it is necessarily the case that $L_t(\omega)$ must reach zero in some finite time as claimed.

For (ii), suppose by way of contradiction that $L_t(\omega_1)$ increases at some time $t \geq \tilde{T}(\omega_1)$. The $q-$complements assumption implies that the marginal value of labor allocated to $\omega_2, \ldots, \omega_{N(s)}$ must increase, while the value of labor allocated to ω_1 decreases. This is because both compute Q and $L_t(\omega_1)$ increase, while the quantity of labor used for $\omega_2, \ldots, \omega_{N(s)}$ decreases (from labor market clearing). This is a contradiction, since the fact that $L_t(\omega_1) > 0$ requires the marginal value of labor to be equalized between ω_1 and the remaining $\omega_2, \ldots, \omega_{N(s)}$.

For (iii), suppose by way of contradiction that $L_t(\omega_1)$ jumps down at time $t \geq \tilde{T}(\omega_1)$. Labor market clearing implies that the quantity of labor used for $\omega_2, \ldots, \omega_{N(s)}$ jumps up. This implies that the marginal value of labor allocated to $\omega_2, \ldots, \omega_{N(s)}$ jumps down (compute affects it continuously), while the marginal value of labor allocated to ω_1 increases (from the q-complements assumption). This generates a contradiction: workers of skill s cannot jump to $\omega_2, \ldots, \omega_{N(s)}$, as this work would pay them a strictly lower wage than ω_1.

In sum, (i)–(iii) show that, after $\tilde{T}(\omega_1)$, employment $L_t(\omega_1)$ decreases gradually and reaches zero in a finite amount of time $\Delta(\omega_1)$. During this transitional period, workers of skill s earn a wage $A_t\ \alpha(\omega_1)$. At time $T(\omega_1)+\Delta(\omega_1)$, we have no more labor in ω_1, which means that, at least until some other automation date is reached, we have a fixed quantity of skill s labor in $\omega_2, \ldots, \omega_{N(s)}$, while Q keeps growing. The $q-$complements assumption implies that from here on and until we reach $\tilde{T}(\omega_2)$, wages will decouple from $A_t\ \alpha(\omega_1)$ and grow while $A_t\ \alpha(\omega_1)$ decreases.

Let's now consider the automation of ω_2. This cannot happen before $\tilde{T}(\omega_1)+\Delta(\omega_1)$. During this period, $W_s(t) = A_t\ \alpha(\omega_1)$. Thus, the only way in which it would make sense to automate ω_2 for the first time in this period is if $\alpha(\omega_1) = \alpha(\omega_2)$, which was ruled out by assumption. After $\tilde{T}(\omega_1)+\Delta(\omega_1)$, wages are growing and exceed $A_t\ \alpha(\omega_1)$. Thus, at the point $\tilde{T}(\omega_2)$ when we automate ω_2, it must be the case that

$$W_t(s) \ = \ A_t\ \alpha(\omega_2)$$

and we necessarily have

$$\alpha(\omega_2) \ > \ \alpha(\omega_1).$$

This shows that, in a transition where compute binds, work in $\Omega(s)$ is automated in order of computational requirements.

We can repeat the same analysis for ω_2, ω_3, and so on, all the way up to $\omega_{N(s)-1}$. The only new detail is in the proof of (iii), where we have to account for the possibility that workers could go back to work that had been previously automated. For example, after $\widetilde{T}(\omega_2)$, one could have labor jumping from ω_2 to ω_1. However, by the time we reach this point, wages are necessarily lower at ω_1 (wages are $A_t\alpha(\omega_1)$) than in ω_2, showing that workers won't reallocate in this direction and the same argument developed above for (iii) applies.

By the time we get to $\omega_{N(s)}$, we have $\alpha(\omega_{N(s)}) = \text{CEU}(s)$, since the most computationally-complex work is saved for last. This work is automated only when $W_t(s)$ reaches a level A_t CEU(s). After this point, wages are pinned down by A_t CEU(s) and all labor of skill s remains employed at $\omega_{N(s)}$. Note that this also describes what would happen if $N(s) = 1$. □

Proof of Proposition 8. Let $T(\omega_1) < \ldots < T(\omega_N)$ be the times at which work becomes automatable. I first show that Y_t jumps up at $T(\omega_i)$. As $t \longrightarrow T(\omega_i)^-$ (from below), we have

$$\underbrace{F_{\omega_i}(\{X_t(\omega)\})}_{V} > \underbrace{A_t \ \alpha(\omega_i)}_{U}.$$

This is because work ω_i is produced with compute the moment this is feasible. Suppose we reach time $T(\omega_i)$. Reallocating a mass $\varepsilon\alpha(\omega_i)$ of compute from other uses to ω_i and keeping all else fixed at its level prior to $T(\omega_i)$ raises output by

$$(V - U)\ \varepsilon > 0.$$

Thus, Y_t must jump up by at least this amount at $T(\omega_i)$.

Suppose now that ω_i is performed by workers of skill s, and that after time $T(\omega)$, only work in $\bar{\Omega}(s) \subset \Omega(s)$ is not yet automatable. Consider three possible cases:

- **Labor still used in ω_i at $T(\omega_i)$ and $\bar{\Omega}(s)$ is non-empty:** I show that, on impact, A_t jumps up, $W_t(s)$ jumps down, and $L_t(\omega_i)$ jumps down.

  At time $T(\omega_i)$, the quantity of labor of skill s and compute used to produce work in $\{\omega_i\} \cup \bar{\Omega}(s)$ jumps up, while the compute used for $\omega_1, \dots \omega_{i-1}$ jumps down. The $q-$ complements assumption implies that the marginal value of compute in $\omega_1, \dots \omega_{i-1}$ (which pins down A_t) jumps up and the marginal value of skill s in $\{\omega_i\} \cup \bar{\Omega}(s)$ jumps down. Therefore, in this case A_t jumps up on impact while $W_t(s)$ jumps down.

  To conclude this case, I show that $L_t(\omega_i)$ necessarily jumps down on impact. Suppose this is not the case, then the quantity of compute and labor used for $\omega_1, \dots \omega_i$ does not change on impact or jumps up. At the same time, the quantity of labor of skill s used for $\bar{\Omega}(s)$ is unchanged or jumps down. In both cases, the $q-$complements assumption implies that the marginal value of skill s in $\bar{\Omega}(s)$ remains unchanged or jumps up. This, however, contradicts the fact that $W_t(s)$ jumps down on impact.

- **Labor fully displaced at $T(\omega_i)$ and $\bar{\Omega}(s)$ is non-empty:** This case necessarily features a jump down in $L_t(\omega_i)$. I show that, on impact, A_t jumps up too, but the effects on $W_t(s)$ are now ambiguous.

At $T(\omega_1)$, the quantity of compute used in $\omega_1, \ldots \omega_i$ is Q and the quantity of skill s used for $\bar{\Omega}(s)$ is $H(s)$. The $q-$complements assumption implies that the marginal value of compute is higher in this allocation than in a counterfactual one where we force $\lim_{t \to T(\omega_i)^-} L_t(\omega)$ workers of skill s to produce ω_i (where they can be pooled with compute) and remove these workers from $\bar{\Omega}(s)$. This is because in such allocation, we would have more resources producing $\omega_1, \ldots \omega_i$ and fewer producing $\bar{\Omega}(s)$. To conclude, note that relative to the initial pre-automation allocation, this counterfactual one features less compute producing $\omega_1, \ldots \omega_{i-1}$ (since some compute would be optimally employed for ω_i), and more resources producing both ω_i and $\bar{\Omega}(s)$. The $q-$complements assumption then implies that the marginal value of compute in the counterfactual allocation exceeds that in the pre-automation one by some positive amount. This chain of reasoning shows that the marginal value of compute jumps up on impact, as claimed.

- $\bar{\Omega}(s)$ **is empty:** Let's now suppose ω_i is the last work in $\Omega(s)$ to be automated. This happens at time $T(s)$. Right before $T(s)$, skill s wages satisfy

$$W_t(s) \; > \; A_t \; \mathrm{CEU}(s).$$

  This is because either $\arg\max_{\omega \in \Omega(s)} \alpha(\omega)$ was automated earlier (and the $q-$complements assumption ensures it remains automated) or because $\omega_i = \arg\max_{\omega \in \Omega(s)} \alpha(\omega)$, in which case the above inequality holds because this work is assumed to be produced with compute the moment this is feasible. At $T(s)$, the maximum wage that workers of skill s can get is $A_t \mathrm{CEU}(s)$, which shows that $W_t(s)$ necessarily jumps down on impact at this time.

To conclude, observe that aggregate wages can be written as $W_t \; = \; Y_t \; - \; A_t Q_t$. Wages jump up if the expansion in output (the productivity effect) dominates the increase in A_t. Otherwise wages may jump down, as shown in Acemoglu and Restrepo (2018). □

Proof of Proposition 9. By assumption, the compute allocated to science is $\lambda \; Q_t$ and the compute allocated to production is $(1 - \lambda) \; Q_t$. The same steps from the proof of Proposition 1 show that all production and scientific work are automated.

The same steps from the proof of Proposition 9 show that, when this happens, output and scientific progress can be written as

$$Y_t \;=\; Z_t\, A\, \big((1-\lambda)\, Q_t \;+\; \sum_{s\in\mathcal{S}} \mathrm{CEU}(s)\, H(s)\big),$$

as in (3), where

$$\frac{\dot{Z}_t}{Z_t} \;=\; Z_t^{-\beta}\, \bar{B}\, \big(\lambda\, Q_t \;+\; \sum_{s\in\mathcal{S}} \mathrm{CEU}(s)\, H(s)\big).$$

Here, $\bar{B} > 0$ is the marginal rate of transformation of compute into innovation.

The unique solution to the differential equation for Z_t is

$$Z_t \;=\; \Big(\beta\, \bar{B}\, \frac{\lambda}{g_Q}\, Q_t \;+\; t\, \beta\, \bar{B} \sum_{s\in\mathcal{S}} \mathrm{CEU}(s)\, H(s) \;+\kappa\Big)^{1/\beta},$$

for some constant κ (pinned down by Z_0). This shows that, as $t \longrightarrow \infty$,

$$\frac{Z_t}{\Big(\beta\, B\, \big(\lambda\, Q_t \;+\; \sum_{s\in\mathcal{S}_R} \mathrm{CEU}(s)\, H(s)\big)\Big)^{1/\beta}} \;=\; \frac{\Big(\beta\, \bar{B}\, \frac{\lambda}{g_Q}\, Q_t \;+\; t\, \beta\, \bar{B} \sum_{s\in\mathcal{S}} \mathrm{CEU}(s)\, H(s) \;+\kappa\Big)^{1/\beta}}{\Big(B\, \big(\lambda\, Q_t \;+\; \sum_{s\in\mathcal{S}_R} \mathrm{CEU}(s)\, H(s)\big)\Big)^{1/\beta}} \longrightarrow 1$$

for $B = \beta\bar{B}/g_Q$, showing the validity of the approximation in (4). □

Proof of Proposition 10. This follows directly from (3) and (4) □

Proof of Proposition 11. Consider the problem of allocating total computational resources between science and production. Let's denote by c_t the share of compute used for production and $\lambda_t = 1 - c_t$ the share used for science. Let's assume c_t and λ_t are bounded away from zero (a point we verify below).

Let's also assume we have reached a time when all bottleneck production and scientific work are automated. From this point onward, the behavior of output and scientific knowledge is

$$Y_t \;=\; Z_t\, A\, \big(c_t\, Q_t \;+\; \underbrace{\sum_{s\in\mathcal{S}} \mathrm{CEU}(s)\, H(s)}_{=H}\big)$$

where

$$\frac{\dot{Z}_t}{Z_t} \;=\; Z_t^{-\beta}\;\bar{B}\;\Big((1-c_t)\;Q_t\;+\;\underbrace{\sum_{s\in\mathcal{S}}\text{CEU}(s)\;H(s)}_{=H_R}\Big) \tag{10}$$

The planner's objective is then to pick $c_t \in (0,1)$ so as to maximize

$$\int_0^\infty e^{-\rho t}\;\frac{(Z_t\;A\;(c_t\;Q_t\;+\;H)^{1-\gamma}}{1-\gamma}\;dt$$

starting from some Z_0 and Q_0, and with production technology evolving according to

$$\dot{Z}_t = Z_t^{1-\beta}\;\bar{B}\;((1-c_t)\;Q_t\;+\;H_R).$$

Throughout, we assume

$$\rho > \Big(1+\frac{1}{\beta}\Big)\;g_Q\;(1-\gamma), \tag{11}$$

which ensures this problem is well define (i.e., has a finite maximum).

The Hamiltonean associated with this problem is

$$H_t(\lambda_t, Z_t, c_t) \;=\; \frac{\big(Z_t\;A\;(c_t\;Q_t\;+\;H)\big)^{1-\gamma}}{1-\gamma}\;+\;\lambda_t\;Z_t^{1-\beta}\;\bar{B}\;((1-c_t)\;Q_t\;+\;H_R),$$

where λ_t is the co-state for Z_t, whose behavior is given by

$$\rho\lambda_t - \dot{\lambda}_t \;=\; \frac{\partial H_t}{\partial Z_t} \;=\; Z_t^{-\gamma}\;\big(A\;(c_t\;Q_t\;+\;H)\big)^{1-\gamma}\;+\;(1-\beta)\;\lambda_t\;Z_t^{-\beta}\;\bar{B}\;((1-c_t)\;Q_t\;+\;H_R).$$

The optimal control condition for c_t satisfies

$$\frac{\partial H_t}{\partial c_t} \;=\; 0 \Leftrightarrow c_t\;+\;\frac{H}{Q_t} \;=\; A^{1/\gamma-1}\;\bar{B}^{-1/\gamma}\;\frac{Z_t^{\frac{\beta}{\gamma}-1}}{Q_t\;\lambda_t^{\frac{1}{\gamma}}}. \tag{12}$$

Plugging in the co-state rule of motion, we get

$$\rho - \frac{\dot{\lambda}_t}{\lambda_t} = m_t \left(c_t + \frac{H}{Q_t}\right) + (1-\beta)\, m_t \left(1 - c_t + \frac{H_R}{Q_t}\right) \tag{13}$$

where $m_t = \bar{B}\, Z_t^{-\beta}\, Q_t$. This measures the effective brainpower of computing resources relative to the difficulty of frontier scientific problems.

Our next goal is deriving a system of equations in c_t and m_t.

For c_t, we can differentiate (12) to obtain:

$$\frac{\dot{c}_t}{c_t + H/Q_t} - \frac{H}{Q_t}\, g_Q$$
$$= \left(\frac{\beta}{\gamma} - 1\right) \underbrace{m_t \left(1 - c_t + \frac{H_R}{Q_t}\right)}_{\text{from } \frac{\dot{Z}_t}{Z_t}} - \underbrace{g_Q}_{\frac{\dot{Q}_t}{Q_t}} - \frac{1}{\gamma} \underbrace{\left(\rho - m_t \left(c_t + \frac{H}{Q_t}\right) - (1-\beta)\, m_t \left(1 - c_t + \frac{H_R}{Q_t}\right)\right)}_{\frac{\dot{\lambda}_t}{\lambda_t} \text{ from (13)}}$$

As $t \to \infty$, $1/Q_t \to 0$ and this simplifies to our first equation

$$\frac{\dot{c}_t}{c_t} = -m_t(1 - c_t) - g_Q - \frac{1}{\gamma}(\rho - m_t). \tag{14}$$

For m_t we obtain

$$\frac{\dot{m}_t}{m_t} = -\beta \underbrace{m_t \left(1 - c_t + \frac{H_R}{Q_t}\right)}_{\text{from } \frac{\dot{Z}_t}{Z_t}} + \underbrace{g_Q}_{\frac{\dot{Q}_t}{Q_t}}.$$

As $t \to \infty$, $1/Q_t \to 0$ and this simplifies to our second equation

$$\frac{\dot{m}_t}{m_t} = -\beta\, m_t\, (1 - c_t) + g_Q. \tag{15}$$

Any solution to the planner's problem must satisfy equations (14), (15), and the transversality condition

$$\lim_{t \to \infty} e^{-\rho t}\, \lambda_t\, Z_t \to 0.$$

The phase diagram for the system formed by (14) and (15) is shown in Figure 3. The figure shows there are three potential types of solution paths.

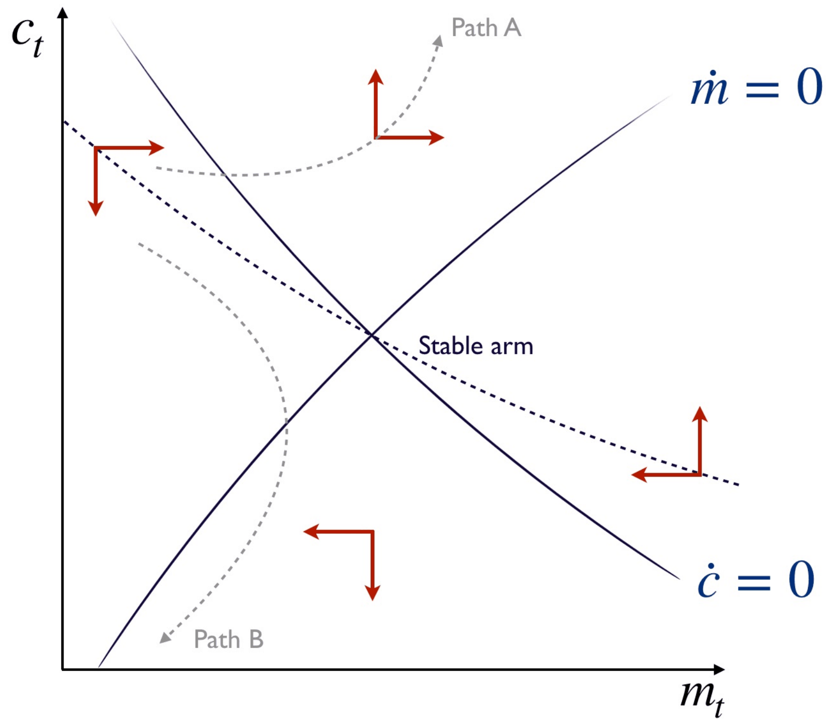

Figure 3: Limit behavior of the system formed by (14) and (15).

First, we have paths converging to the northeast, labeled by the letter A in the figure. These paths become infeasible, since they involve $c_t > 1$.

Second, we have paths converging to $c_t = 0$ and $m_t = g_Q/\beta$, labeled by the letter B in the figure. These paths violate transversality, since the growth rate g_T^B of $e^{-\rho t}\ \lambda_t\ Z_t$ along this path is

$$g_T^B \ = -\rho \ + \underbrace{\rho \ - \ (\frac{1}{\beta} - 1)\ g_Q}_{\text{from } \frac{\dot{\lambda}_t}{\lambda_t}} \ + \underbrace{\frac{g_Q}{\beta}}_{\text{from } \frac{\dot{Z}_t}{Z_t}} \ = \ g_Q > 0.$$

Finally, we have a unique path along the stable manifold converging to a steady state, given by

$$\begin{aligned} c^* &= 1 - \frac{g_Q}{\beta\rho + (1+\beta)\gamma g_Q}, \\ m^* &= \rho + \gamma g_Q (1 + \frac{1}{\beta}). \end{aligned}$$

Along this path, the transversality condition holds. This is because the growth rate g_T^* of $e^{-\rho t}\ \lambda_t\ Z_t$

along this path is

$$g_T^* \;=\; -\rho \;+\; \underbrace{\rho \;-\; m^*\, c^* - \;(1-\beta)\, m^*\, (1-c^*)}_{\text{from } \frac{\dot\lambda_t}{\lambda_t}} \;+\; \underbrace{\frac{g_Q}{\beta}}_{\text{from } \frac{\dot Z_t}{Z_t}} \;.$$

Substituting m^* and c^*, this growth rate becomes

$$g_T^* \;=\; -\rho + \left(1 + \frac{1}{\beta}\right)\, g_Q\, (1-\gamma),$$

which is negative, from (11), establishing the transversality condition.

The above argument shows that the only possible solution to the planners' problem lies along the stable manifold. Because there is a solution to the problem, it is the one along this path. Moreover, along this path, λ_t converges to the long-run value for λ^* in the proposition. Note that inequality (11) ensures $c^* \in (\beta/(1+\beta), 1)$, or equivalently, $\lambda^* \in (0, 1/(1+\beta))$. This shows that, as initially guessed, the quantity of compute assigned to both uses grows proportionally along the optimum.

To conclude, note that starting from any $m_0 > m^*$ (i.e., a point where the level of scientific knowledge is low, relative to available computing resources), the solution to the planner's problem features an increasing path for c_t, or equivalently, a decreasing path for λ_t, as claimed in the proposition. □