When Does Automating AI Research Produce Explosive Growth?

Feedback Loops in Innovation Networks

Tom Davidson[*]   Basil Halperin[†]   Thomas Houlden[‡]   Anton Korinek[§]

January 2026

Abstract

AI labs are increasingly using AI itself to accelerate AI research, creating a feedback loop that could potentially lead to an "intelligence explosion". We develop a general semi-endogenous growth model with an innovation network, where research and automation in one sector increases the productivity of research in other sectors, and derive a clean analytical condition under which growth becomes superexponential ("explosive"). The key intuition is that automation of research both offsets diminishing returns to research and increases cross-sectoral research spillovers, making explosive growth more likely. Applying this model to a calibrated, AI-integrated economy, we demonstrate that the growth effects of automation may be slow initially but compound rapidly. In our benchmark calibration, the level of automation needed to double the long-run growth rate already achieves well over half of the automation level needed to generate explosive growth.

1 Introduction

> "Advanced AI is interesting for many reasons, but perhaps nothing is quite as significant as the fact that we can use it to do faster AI research." (Altman, 2025)

AI has the potential to automate many different kinds of work: customer support, coding, marketing, and many other tasks. However, a widespread belief among leading AI researchers is that automation of *AI research itself* in particular will have a transformative economic impact. Central to this thesis is the argument that such *recursive self-improvement* – where AI systems become increasingly capable of designing and improving themselves – creates a feedback loop leading to an "intelligence explosion" and rapid economic growth (Good, 1966, Yudkowsky, 2013). OpenAI bluntly declares its goal of developing such technology "by March of 2028" (OpenAI, 2025).

Economists have traditionally been skeptical about the possibility of explosive growth from recursive self-improvement, pointing to two obstacles:

1. **Diminishing returns.** A self-improving process *may* achieve hyperbolic growth ("explode") – but, contra Good (1966), such a process does not *necessarily* explode. Whether or not a recursively self-improving process explodes depends critically on the strength of *diminishing returns*. Formally, a process with a positive feedback loop $\frac{dy}{dt} = y^{1-\beta}$ will not explode if there are diminishing returns, i.e. if $\beta \geq 0$. Intuitively, a self-improving AI may "pick all the low hanging fruit first" and find it increasingly difficult to make algorithmic progress, and as a result progress may be subexponential or even stagnate. In the economics literature, this corresponds to models where "ideas get harder to find" (Bloom et al., 2020).

2. **Bottlenecks.** Even if one process in the economy achieves explosive growth, this does not necessitate that *aggregate* growth explodes: progress in one sector may be *bottlenecked* by slow progress in other sectors (Aghion, Jones and Jones, 2019, Jones, 2025, Jones and Tonetti, 2025). Intuitively, even if AI was capable of producing infinite left shoes, "total shoe output" would still be bottlenecked by production of right shoes. Formally, this is captured by *complementarity*: in its most extreme form, total output is the minimum of two components, $y = \min\{y_1, y_2\}$.

This paper. In this paper, we set aside the issue of bottlenecks and study the effect of automating AI R&D on economic growth by focusing on how such automation will *offset diminishing returns* in research – potentially even eliminating diminishing returns – by creating and amplifying *feedback loops*.

We begin by writing down a fully general extension to the canonical semi-endogenous growth model, which may be of independent interest to the growth literature away from AI. The model has three key features, each of which is applicable to the economy more generally but is of particular importance in our setting.

(1) *"Technological* feedback loops" across an innovation network. In the canonical model, there is one research sector. Our model features a network of heterogeneous research sectors, where innovations in one sector spill over to increase the rate of innovation in other sectors (Liu and Ma, 2024, Ngai and Samaniego, 2011). Such spillovers are important for capturing, for example, the feedback loop between better software and better hardware: better computer chips allow for OpenAI to design better AI models, while better AI models in turn are used to help to design yet better computer chips (Mirhoseini et al., 2021), and so on.

(2) *"Economic* feedback loops". Economic feedback loops refer to the channel where higher output is transformed back into a driver of further economic growth. The classic example is capital accumulation: higher output leads to more savings and investment, which increases the capital stock, which produces yet more output, and so on.

In our setting, economic feedback loops are particularly important in how they *interact* with technological feedback loops. This captures the idea that AI-induced technological progress increases aggregate GDP, which is necessary to fund further AI R&D investment. For example, the all-in cost of building frontier AI models has grown roughly by a factor of 10 every two years continuously for the last six years (Cottier et al., 2024, Whitfill, Snodin and Becker, 2025, Nesov, 2025). If this rate were to continue, the cost of building one individual frontier model in 2030 would be in excess of $1 trillion. Such investment growth seems unlikely to continue – unless AI progress itself can raise total output and therefore the quantity of available resources for AI investment.

(3) Automation. We introduce the idea of automation of the ideas production function (Aghion, Jones and Jones, 2019) into a network context, demonstrating how automation in one research sector spills through the innovation network.

Analytical insights. The general model produces simple analytical conditions under which technological and economic feedback loops give rise to balanced or explosive growth. It also provides several broad, interpretable insights.

- Technological feedback loops – spillovers across research sectors – directly offset diminishing returns. In other words, they reduce the degree to which ideas get harder to find. If such spillovers are improperly ignored, a system may be estimated to be non-explosive when in reality spillovers from other sectors may tip the system into explosive dynamics. More prosaically, the recognition of the existence of the innovation network affects estimation of canonical ideas-getting-harder-to-find parameters by sector.

- Automation creates economic feedback loops, therefore effectively offsetting diminishing returns. Automation means that a task which was previously performed by human labor is instead performed by machines, i.e. capital. Human labor does not have an economic feedback loop in the modern era: higher GDP does not result in a higher population. On the other hand, machines have an economic feedback loop: higher GDP results in the construction of more machines. Thus, replacing human labor with capital offsets diminishing returns by creating a new feedback loop.

- The interaction between technological and economic feedback loops amplifies each. This is the idea that AI research results in higher GDP, which helps fund further investment in AI research.

Calibrated application: AI automation with software & hardware feedback loops.
We apply the general framework to study an AI economy, modeled to match key features of modern AI development, to analyze the titular question of how automation of AI research affects economic growth. While the model is rich and complex, it can be studied analytically, and is summarized in figure 1.

The starting assumption of this model is that AI can automate some fraction of tasks across different sectors in the economy. AI itself is a combination of a nonrivalrous idea, "software" (AI algorithms), and a kind of capital, "hardware" (computer chips like Nvidia GPUs). Software progress follows a canonical ideas production function. Computer chip hardware, meanwhile, accumulates like any form of capital, augmented by investment-specific technical change: "hardware quality" is another standard nonrivalrous idea following an ideas production function. In addition to software and hardware quality, we have a third innovation sector: a "general" research sector creating new ideas for goods production, as in standard models. Production of new ideas in any of the three innovation sectors may be performed by some combination of human labor

3

or automated by AI. Finally, goods output likewise may be produced by a combination of humans or automated by AI.

We then use the insights of this model to illustrate how automation of both research and production can amplify existing feedback loops or spawn them where they did not exist before. In particular, the model predicts hyperbolic growth will arise under the simple condition:[1]

$$f_Y + f_S \left(\frac{1}{\beta_S} \right) + f_H \left(\frac{1}{\beta_H} \right) + f_A \left(\frac{1}{\alpha} \right) \left(\frac{1}{\beta_A} \right) > 1 \qquad (1)$$

Here, the subscripts denote sectors: output Y, software S, hardware H, and general innovation A. The term $f_i \in [0,1]$ is the fraction of tasks in each sector that can be automated by AI. The term $\beta_i > 0$ is the degree of diminishing returns in each research sector.[2]

Thus, the explosion condition (1) says that the economy features hyperbolic growth if the sum of the strength of feedback loops is greater than unity. In particular, the relevant feedback loops are: (1) the pure economic feedback loop (which will be strengthened with f_Y); (2) the feedback loops induced through technological and economic channels as automation accelerates research (the next three terms, f_i scaled down by diminishing returns $1/\beta_i$).[3]

We emphasize that this condition does not pin down an exact growth path or the timing of a growth explosion. Rather, it specifies the condition that determines whether growth will *eventually* explode. Hence, we think of this condition as a line in the sand: supposing other parameters are fixed, how much automation f_i must be achieved to tip the economy into an explosive growth regime?

Empirical insights. The explosion condition highlights the critical importance of measurement of the degree of diminishing returns in the software and hardware sectors, β_H and β_S. Bloom et al. (2020) estimate that in the economy as a whole, ideas become sharply harder to find, with $\beta_A = 3.1$. In comparison, while the hardware sector does feature the same phenomenon – more and more researchers are required to maintain the pace of Moore's Law – the quantitative magnitude is *much* smaller: $\beta_H = 0.2$. Indeed, the hardware sector shows the smallest degree of these diminishing returns of

---

[1]For simplicity, the expression here assumes no "parallelization penalty"; (54) generalizes.
[2]α is the labor share in production of output.
[3]In the case of general innovation, A, diminishing returns are modulated by the labor share, α.
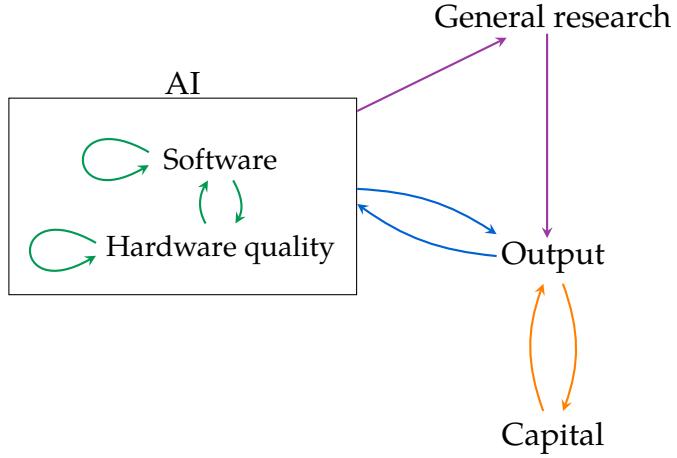
Figure 1: Model of AI automation with software & hardware feedback loops

any sector studied!

Estimating the degree of diminishing returns in software research has proven more challenging and is an important avenue for future research. The best evidence, from Ho and Whitfill (2025) and Erdil, Besiroglu and Ho (2024), estimates $\beta_S \approx 1$. With this calibration, we can estimate the extent of automation needed to achieve hyperbolic growth in this model.

Limitations. As noted, these results are derived under the strong assumption of an elasticity of substitution of unity across tasks (Cobb-Douglas aggregation) to turn off the issue of bottlenecks. Under this aggregation, output can be driven by explosive progress in a subset of tasks. This means that automation of some tasks is enough to kickstart explosive dynamics across the system. The hyperbolic growth threshold we arrive at is strikingly low. This suggests that one of the most important areas for additional research is understanding the degree of substitutability between economic tasks in a world with increasing automation. We discuss additional limitations of the model developed here in Section 6.

Related literature. Several studies have investigated the possibility of AI and automation leading to transformative growth. Aghion, Jones and Jones (2019) characterize the conditions under which a single-sector system capable of recursive self-improvement exhibits hyperbolic growth. Trammell and Korinek (2025) extend this analysis by embedding such self-improving technologies in a macroeconomic setting where capital

accumulation and technological progress reinforce each other. Jones (2025) quantitatively studies the role of bottlenecks in preventing explosive growth and Jones and Tonetti (2025) estimate the size of these bottlenecks historically.

We advance this line of work in three main directions. First, we generalize the analysis to systems whose progress depends on a network of multiple technologies with heterogeneous returns to research effort, providing a unified condition for when growth becomes explosive. Second, we integrate technological heterogeneity with economic feedback loops, showing how output–technology complementarities amplify or dampen the possibility of hyperbolic growth. Third, we use this richer framework to derive a threshold for automation that marks the transition from balanced to hyperbolic growth and provide quantitative estimates under empirically grounded parameters.

Erdil et al. (2025) make a uniquely rich effort to model the growth implications of AI using an Integrated Assessment Model. Our simpler framework captures the same core dynamics—with AI displacing human labor in both research and production, thereby driving growth and accelerating further AI progress—while remaining analytically tractable. This tractability allows us to characterize the interaction between economic and technological feedback loops and to derive conditions under which growth becomes explosive.

Networked models of technological progress have also been applied to understand optimal R&D allocations (Liu and Ma, 2024), as well as the sources of heterogeneity in sectoral productivity growth (Ngai and Samaniego, 2011). These analyses explicitly rule out explosive growth by imposing constant returns to scale on cross-sector innovation spillovers, which ensures balanced growth in an endogenous growth setting. We introduce automation into such a model, apply it to the particular setting of software- and hardware-driven AI progress, as well as study the possibility of explosive growth.

Outline. This paper proceeds as follows. Section 2 develops a set of simple models to illustrate the core economic forces at work in our setting. Section 3 presents the general model of semi-endogenous growth with an innovation network. Section 4 introduces AI-driven automation into the general model. Section 5 presents our simple integrated AI economy, applying the results from the general model. Section 6 discusses before concluding.

2   A sequence of simple models

As is typical in networked settings, our general model is fairly complicated. This section presents a sequence of simple models to highlight the core economic forces at work in the general model. We draw out four lessons:

1. Diminishing returns prevent growth explosions.
2. Innovation networks (*technological* feedback loops) introduce spillovers and offset diminishing returns.
3. Economic feedback loops also introduce spillovers and offset diminishing returns.
4. Automation introduces new spillovers (or strengthens existing ones).

2.1   Lesson one: Diminishing returns prevent growth explosions

"Let an ultraintelligent machine be defined as a machine that can far surpass all the intellectual activities of any man however clever. Since the design of machines is one of these intellectual activities, an ultraintelligent machine could design even better machines; there would then unquestionably be an 'intelligence explosion', and the intelligence of man would be left far behind."
— Good (1966), Speculations Concerning the First Ultraintelligent Machine

The simplest possible formalism for I.J. Good's concept of an intelligence explosion is $\dot{S}_t = S_t$, where S is the level of intelligence or the level of "software productivity", and dots indicate time derivatives. This equation says that when the *level* of intelligence S_t is low, the *rate of change* of intelligence $\dot{S}_t$ is also low; and when the level of intelligence is high, the rate of change of intelligence is also high.

We can generalize this process:

$$\dot{S}_t = S_t^{1-\beta} \tag{2}$$

With this generalization, we can observe that – contra Good (1966) and many since – there need not be an "explosion" from a recursively self-improving process. In particular, if $\beta < 0$, so that there are increasing returns, then there is a literal mathematical singularity: S_t approaches infinity in finite time. On the other hand, if $\beta = 0$, the process exhibits exponential growth, and if $\beta > 0$, so that there are diminishing returns,

the process is subexponential or even sublinear. This is a simple reminder of the importance of diminishing returns in preventing runaway feedback processes.

Equation 2 is in fact the canonical form of the ideas production function for modeling the growth of productivity (abstracting from research inputs for now), and it can be easily embedded in an economic growth model to think about the relationship between intelligence explosions and economic explosions. The simplest possible case has a goods production function as follows, assuming an exogenous bounded path for the supply of labor L_t:

$$Y_t = S_t L_t^{\alpha} \tag{3}$$

Output is produced with labor input (subject to potentially diminishing returns, $\alpha \in [0, 1]$) augmented by software capabilities.

The simple economy of (2)-(3) clearly features an *economic* explosion – infinite output Y in finite time – if and only if there is an intelligence explosion, i.e.,

$$\beta < 0 \tag{4}$$

This model is summarized in figure 2.
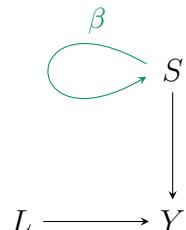
Figure 2: Recursive self-improvement explodes if and only if there are no diminishing returns: $\beta < 0$.

2.2 Lesson two: Innovation networks introduce spillovers and offset diminishing returns

It is well-known that higher quality computer chips are used by AI researchers to write better algorithms; it is additionally, increasingly the case that those better AI algorithms are used in turn to help design better chips. "AlphaChip" (Mirhoseini et al., 2021)

from Google DeepMind is a striking example of this phenomenon. A reinforcement learning method for designing chip layouts, AlphaChip is reported to have been used in designing every new generation of Google's Tensor Processing Unit chip since 2020 and to be responsible for a growing share of the 'floorplan' for each generation of chip (Goldie and Mirhoseini, 2024).

To capture this, we modify the baseline semi-endogenous growth model of (2)-(3) to introduce a two-sector, *networked* semi-endogenous growth model. Continue to denote S as software productivity, denoting H as hardware quality, and dropping time subscripts to ease notation,[4]

$$\dot{S} = S^{1-\beta_S} H^{\phi_S} \tag{5}$$

$$\dot{H} = H^{1-\beta_H} S^{\phi_H} \tag{6}$$

$$Y = (SH)^{1/2} L^{\alpha} \tag{7}$$

Here, β_S now denotes diminishing returns within software ("ideas getting harder to find"); likewise, β_H the same within hardware. Meanwhile $\phi_S \geq 0$ reflects the *technological spillovers* from hardware quality to software improvements, and vice versa for ϕ_H. This nests the dynamics of the prior model under $\phi_S = \phi_H = 0, \beta_H = 1$.
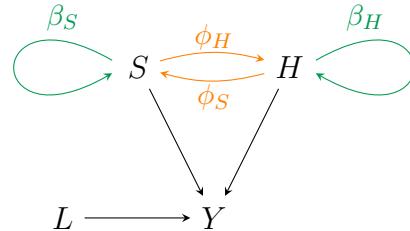
Figure 3: Innovation networks introduce spillovers, offsetting diminishing returns.

In this networked example, we now have three feedback loops, which can be seen by physically tracing all possible "loops" in figure 3:

1. Recursive self-improvement within software, as before, governed by β_S.

2. Recursive self-improvement within hardware quality, governed by β_H.

3. Spillovers across the innovation network, intermediated via ϕ_H and ϕ_S.

---

[4]The choice of equal-weighted Cobb-Douglas aggregation of S and H in the goods production function (7) is not essential.

The spillovers across the innovation network are summarized by the interaction matrix: the matrix collecting the exponents in (5)-(6):

$$\begin{bmatrix} 1 - \beta_S & \phi_S \\ \phi_H & 1 - \beta_H \end{bmatrix}$$

The system can now explode in two ways. First, analogously to the prior single-sector example, the system explodes if either recursive self-improvement loop is strong enough on its own, $\beta_S < 0$ or $\beta_H < 0$.

Second, the system explodes via the spillover loop *if the interaction between the two loops is strong enough*. It turns out that, mathematically, this occurs if the interaction matrix has an eigenvalue greater than unity. This generalizes the single-sector condition that the exponent in the law of motion, $1 - \beta$, is greater than unity. In turn, it can be shown that the eigenvalue condition holds here if and only if:

$$\underbrace{\beta_S \cdot \beta_H}_{\substack{\text{diminishing} \\ \text{returns}}} < \underbrace{\phi_S \cdot \phi_H}_{\text{spillovers}} \tag{8}$$

Condition (8) implies that spillovers effectively offset diminishing returns. For example, suppose $\beta_H = 1$, so that the only difference with the model in section 2.1 is the spillovers. Then the condition (8) for explosive growth becomes simply $\beta_S < \phi_S \phi_H$. Thus, explosive growth no longer requires increasing returns, $\beta_S < 0$, but now can occur if diminishing returns are mild, $\beta_S \in [0, \phi_S \phi_H)$.

2.3 Lesson three: Economic feedback loops introduce spillovers and offset diminishing returns

In our full AI economy model, we combine the technological feedback loops from an innovation network of section 2.2 with *economic* feedback loops. An "economic" feedback loop refers to a feedback loop when higher output is involved.

The most basic economic feedback loop is a Solow model without technological progress: normalizing the population to one, output is produced as $Y = K^{1-\alpha}$ and capital accumulates as $\dot{K} = aY - \delta K$, where a is a constant savings rate and δ the depreciation rate. Of course, this model features explosive growth if there are increasing returns to capital in production $-\alpha < 0$ – or in the language used here, if the economic

10

feedback loop is sufficiently strong.

In the rest of this subsection, we illustrate an economic feedback loop using the canonical single-sector semi-endogenous growth model, with capital instead of labor in the ideas production function:

$$\dot{A} = A^{1-\beta} \cdot (\kappa_A K)^{1-\gamma} \tag{9}$$

$$Y = A \cdot L^{\alpha} \cdot (\kappa_Y K)^{1-\alpha} \tag{10}$$

$$\dot{K} = aY - \delta K \tag{11}$$

Here, A is a general productivity term produced from general research in the economy, replacing the previous S term. (κ_A is the share of capital used for research; $\kappa_Y \equiv 1 - \kappa_A$ is the share used in production.) This model is summarized in figure 4.
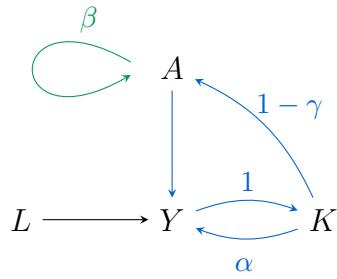

Figure 4: Economic feedback loops effectively offset diminishing returns.

The same logic can be applied as before, with the same condition on the eigenvalues of the interaction matrix. The condition implies the system explodes if:

$$\underbrace{\beta \cdot \alpha}_{\substack{\text{diminishing} \\ \text{returns}}} < \underbrace{(1-\gamma) \cdot 1}_{\text{spillovers}} \tag{12}$$

This condition is exactly analogous to (8) and highlights that when higher productivity A increases output, then if this output in turn can be invested to produce yet further research advances ($\gamma < 1$), then explosive growth is more likely.

Notably, a standard calibration of (12) would imply a lack of explosive growth. Using $\beta = 3.1$ as the extent to which ideas are getting harder to find in the economy as a whole (Bloom et al., 2020), $\alpha = 0.6$ as the labor share in production, and $1 - \gamma = 0.1$ as the capital share in R&D for the economy as a whole (Besiroglu, Emery-Xu and Thomp-

11

son, 2024), we find that the explosion condition is far from being met. However, as the next section shows, *increasing automation* could change this.

2.4 Lesson four: Automation introduces new spillovers

Finally, we come to the role of automation, which is critical to our titular question. We consider automation in a task-style framework, where as the simplest example tasks X_i are bundled into aggregate output Y via a Cobb-Douglas aggregate:[5]

$$Y = \prod_{i=1}^{N} X_i^{1/N}$$

Individual tasks can be produced either with capital or with labor:

$$X_i = \begin{cases} L_i & \text{if not automated} \\ K_i & \text{if automated} \end{cases}$$

Suppose only tasks $i = 1, ..., I$ are automated by capital. Then, optimally spreading inputs equally across tasks, we can write an *effective* aggregate production function:

$$Y = L^{1-f} K^f \xi_Y$$

where ξ_Y is an unimportant constant, and importantly f is defined to measure the share of automated tasks:

$$f \equiv I/N$$

As a result, for our purpose of studying explosive dynamics, automation of tasks can be understood as *increasing the capital share f*: shifting production weight from L to K.

Automation of the *ideas* production function can likewise be microfounded, after adding labor as a factor of production. We will use $f_Y \in [0, 1]$ to denote the share of automated tasks in goods production and $f_A \in [0, 1]$ for the share of automated tasks in ideas production.

Thus, we can simply take our previous system (9)-(15), and consider changes in the

---

[5]This task aggregator rules out bottlenecks by imposing an elasticity of substitution of one across tasks; section 6 discusses this important assumption.

capital share in both production functions:

$$\dot{A} = A^{1-\beta} \cdot (\ell_A L)^{\gamma(1-f_A)} \cdot (\kappa_A K)^{(1-\gamma)+f_A\gamma} \xi_A \tag{13}$$

$$Y = A \cdot (\ell_Y L)^{\alpha(1-f_Y)} \cdot (\kappa_Y K)^{(1-\alpha)+f_Y\alpha} \xi_Y \tag{14}$$

$$\dot{K} = aY - \delta K \tag{15}$$

Here, $f_A \in [0, 1]$ is the degree of automation in the software sector, and $f_Y \in [0, 1]$ is the degree of automation in goods production. (ℓ_A is the share of labor used for research; $\ell_Y \equiv 1 - \ell_A$ is the share used in production; ξ_Y and ξ_A are unimportant constants.)

This system, visualized in figure 5, make clear how automation either strengthens existing feedback loops – or creates new ones which did not exist previously.
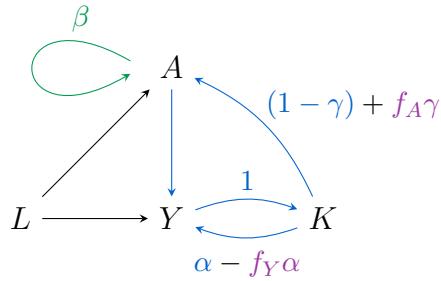

Figure 5: Automation introduces new feedback loops or strengthens existing ones.

For example, suppose initially capital was not used at all in production of ideas and no relevant tasks were automated: $\gamma = 1$ and $f_A = 0$. Then the system would have no economic feedback loops: there would be no arrow from K to A, breaking the loop. If automation f_A then begins creeping above zero, this creates an economic feedback loop.

Alternatively, the effect of automation can be interpreted as directly offsetting diminishing returns (since spillovers offset diminishing returns). This can be seen on the diagram as automation raising the strength of various edges.

The formal condition for a growth explosion is once again exactly analogous to the previous condition, equation (12), simply with automation-augmented terms:

$$\underbrace{\beta \cdot \alpha(1 - f_Y)}_{\text{diminishing returns}} < \underbrace{(1 - \gamma + f_A\gamma) \cdot 1}_{\text{spillovers}} \tag{16}$$

Automation of output offsets diminishing returns to capital accumulation; automation

of ideas here increases spillovers through the economic feedback loop. In section 4, we will see that automation *which is AI-induced* also creates spillovers through a *technological* feedback loop.

2.5 Summary

We can now return to the condition of the introduction, (1), to provide some intuition for its origin. Recall that condition:

$$
f_Y + f_S \left(\frac{1}{\beta_S} \right) + f_H \left(\frac{1}{\beta_H} \right) + f_A \left(\frac{1}{\alpha} \right) \left(\frac{1}{\beta_A} \right) > 1
$$

Now consider the condition we derived in lesson 4, (16). Our main application sets the initial capital share in research of zero, which would be equivalent to setting $\gamma = 1$.[6] Using this, the condition (16) can easily be rewritten as:

$$
f_Y + f_A \left(\frac{1}{\alpha} \right) \left(\frac{1}{\beta} \right) > 1
$$

Clearly, this condition matches the first and last terms of the boxed condition (as the "β" of this formula maps to "β_A" in the former). The two missing terms will come from incorporating a software sector, incorporating a hardware research sector, and introducing a notion of *AI-driven* automation.

The rest of the paper. The sequence of four simple models above illustrates the core economic forces at work. They also show a striking degree of formal mathematical parallels. We now turn to a general framework that explains the deeper underlying structure.

3 A general framework for hyperbolic growth

In this section, we present a general framework to think about the ideas introduced in section 2. We begin in 3.1 by introducing a general innovation network, with spillovers generating arbitrary possible feedback loops between technologies, to consider the necessary and sufficient conditions for hyperbolic growth from technological feedback

---

[6] Setting $\gamma < 1$ would only increase the likelihood of a growth explosion.

loops alone. We then embed this model of networked technological progress into an economic environment in 3.2, demonstrating economic feedback loops that can be isolated analytically in their contribution balanced or explosive growth. Section 4 introduces AI-driven automation.

3.1 Technological feedback loops

Consider an economy with N different technological sectors. Progress in any one sector, $i \in I$, benefits from spillovers from other sectors:

$$\dot{A}_i = v_i R_i^{\lambda_i} \prod_{j \in I} A_j^{\phi_{i,j}} \tag{17}$$

where A_i is the level of technology in sector i, $R_i = L_i^{\gamma_i} K_i^{1-\gamma_i}$ is the aggregated capital and labor research inputs to the sector, and v_i is a constant scaling parameter. Here, since v_i is the only constant variable and is ultimately unimportant for analysis, we drop time subscripts entirely and note that all capitalized variables are growing over time. There are intratemporal diminishing returns to parallel research input, $\lambda_i \in (0, 1]$, and sectoral spillovers, $\phi_{i,j} \geq 0$. We define $\phi_{i,i} = 1 - \beta_i$, where β_i captures the degree of intertemporal diminishing returns to research within a given sector as introduced in Jones (1995). We impose $\beta_i > 0$ so ideas are getting harder to find; otherwise the system necessarily generates explosive growth.[7]

A balanced growth path for this innovation network occurs if all technologies grow at a constant rate, $\dot{A}_i / A_i = g_{A_i}^{\text{BGP}}$ constant. From equation (17), we can see that *if* a balanced growth path exists, then:

$$g_{A_i}^{\text{BGP}} = \frac{\lambda_i}{\beta_i} g_{R_i} + \sum_{j \in I \backslash i} \frac{\phi_{i,j}}{\beta_i} g_{A_j}^{\text{BGP}} \tag{18}$$

That is, the growth rate of technology i on a BGP equals the growth rate of research inputs, g_{R_i}, plus a spillover-weighted sum of the growth rate in other sectors j; all scaled down by the degree to which ideas get harder to find, β_i.

It will be useful to define labels for two of the terms of (18), as they will appear

---

[7]Although, if $N = 1$, $\beta_i = 0$ would be insufficient for explosive growth (Aghion, Jones and Jones, 2019). However, we are generally interested in multi-technology systems where any one sector with $\beta_i = 0$ would be sufficient for explosive growth.

repeatedly in our analysis:

$$r_i = \frac{\lambda_i}{\beta_i}$$

$$s_{i,j} = \frac{\phi_{i,j}}{\beta_i} \qquad \text{for } i \neq j \tag{19}$$

The term r_i is a sector-specific measure of research productivity, capturing how technological progress responds to 'own-sector' research effort. We also introduce an analogous term spillover term, $s_{i,j}$ capturing how technological progress in sector j impacts technological progress in sector i.

We can simplify this system further by writing it in matrix form. Define the $N \times N$ technological feedback matrix, $\mathbf{F}^A$, as:

$$\mathbf{F}^A_{i,j} = \begin{cases} s_{i,j}, & \text{if } i \neq j \\ 0, & \text{if } i = j \end{cases}. \tag{20}$$

From here we state the central mathematical result which we will subsequently apply to a variety of growth environments:

Proposition 1. *Take a (strictly positive) growth system defined by equations for each $i \in I$*

$$\dot{A}_i = v_i E_i^{\ell_i} A_i^{1-b_i} \prod_{j \in I/i} A_j^{p_{i,j}}$$

with $b > 0$ and $p \geq 0$, and the (irreducible) matrix $\mathbf{F} \in \mathbb{R}^{N \times N}_{\geq 0}$ where $\mathbf{F}_{i,j} = p_{i,j}/b_i$ for $i \neq j$, $\mathbf{F}_{i,i} = 0$ and variables E grow at a constant, exogenous rate $g_E \in \mathbb{R}^N_{\geq 0}$. The spectral radius of $\mathbf{F}$, $\rho(\mathbf{F})$, partitions the growth system into three cases:

1. *$\rho(\mathbf{F}) < 1$. The system exhibits balanced growth along the path*

$$\boxed{g_A^{BGP} = \mathbf{\Psi} \mathbf{r} g_E} \tag{21}$$

   where $\mathbf{\Psi} = (\mathbf{I} - \mathbf{F})^{-1} \in \mathbb{R}^{N \times N}_{\geq 0}$ and $\mathbf{r} = diag(\ell_i/b_i, \dots, \ell_N/b_N)$.

2. *$\rho(\mathbf{F}) = 1$. The system growth at any time is bounded by double-exponential growth, and exponential in the purely endogenous growth case ($g_E = 0$).*

3. *$\rho(\mathbf{F}) > 1$. The system exhibits hyperbolic growth, with all variables growing to infinity in finite time.*

16

Proof. See Appendix A.   □

This result implies that to understand the conditions that give rise to explosive growth, we can generally limit our focus to the behavior of the balanced growth path – or more specifically, whether it exists. In our spillover model, the feedback matrix $\mathbf{F}^A$ conveniently provides (i) the balanced growth path via the (technological) Leontief inverse, $\mathbf{\Psi_A} = (\mathbf{I} - \mathbf{F}^A)^{-1}$ and (ii) the conditions where the system in equation (17) exhibits hyperbolic growth. $\mathbf{\Psi}_A$ captures how progress in each sector feeds back into all other sectors repeatedly along the balanced growth path; while r captures the direct effect of research inputs within a sector. We can see that entries in $\mathbf{\Psi}_A$ will be increasing with returns to other sector research, $s_{i,j}$, and therefore increasing with spillover terms $\phi_{i,j}$ and decreasing with the strength of diminishing returns to research β_i.

This balanced growth path nests that of the standard semi-endogenous balanced growth path presented in Jones (1995); setting $\phi_{i,j} = 0$ (so $\mathbf{\Psi}_A = \mathbf{I}$) for $i \neq j$ and allowing research inputs to be proportional to population growth we have the vector of balanced growth paths

$$g_A^{\text{BGP}} = \mathbf{r}n.$$

Liu and Ma (2024) also note that explosive growth can be inferred directly from the eigenvalues of the exponent matrix (of ϕ terms) in equation (17). However, the balanced-growth specification is particularly useful here since empirical estimations of returns to research effort that we ultimately use to calibrate a threshold from explosive growth from automation are calculated under the assumption of a balanced growth path.

In Figure 6 we illustrate hyperbolic and balanced growth in phase diagrams of a two technology model. From this figure we can see stability of the balanced growth path—where $\hat{g}_A$, the growth in the growth rate of A, is zero—can be understood as a crossing condition on the isolines. In this case, the slope of the $\hat{g}_{A_1}$ isoline is $s_{1,2}$ and is $s_{1,2}^{-1}$ for the $\hat{g}_{A_2}$ isoline.[8] Therefore, lines cross whenever $s_{1,2}s_{2,1} > 1$. This is the exact case where the largest eigenvalue of $\mathbf{F}^A$ exceeds one and there is no well-defined $\mathbf{\Psi}_A$.

Intuitively, in the two-sector example, the explosive growth condition $s_{1,2}s_{2,1} > 1$ is more likely when either feedback term $s_{1,2}$ or $s_{2,1}$ is large. Recall that the feedback terms are defined as $s_{i,j} = \phi_{i,j}/\beta_i$, where $\phi_{i,j}$ measures the spillover of sector j to growth in

---

[8]One can arrive at this result by taking the time derivative of equation 17 and then setting the growth-in-growth rates, $\hat{g}_{A_1}$ and $\hat{g}_{A_2}$, to zero and rearranging λ, ϕ and β terms into s and r terms by their definitions in equation 19.

Figure 6: Hyperbolic vs Balanced Growth in a Two Technology Network

(a) Stable, Balanced Growth ($s_{2,1}^{-1} > s_{1,2}$)

(b) Hyperbolic Growth ($s_{2,1}^{-1} < s_{1,2}$)
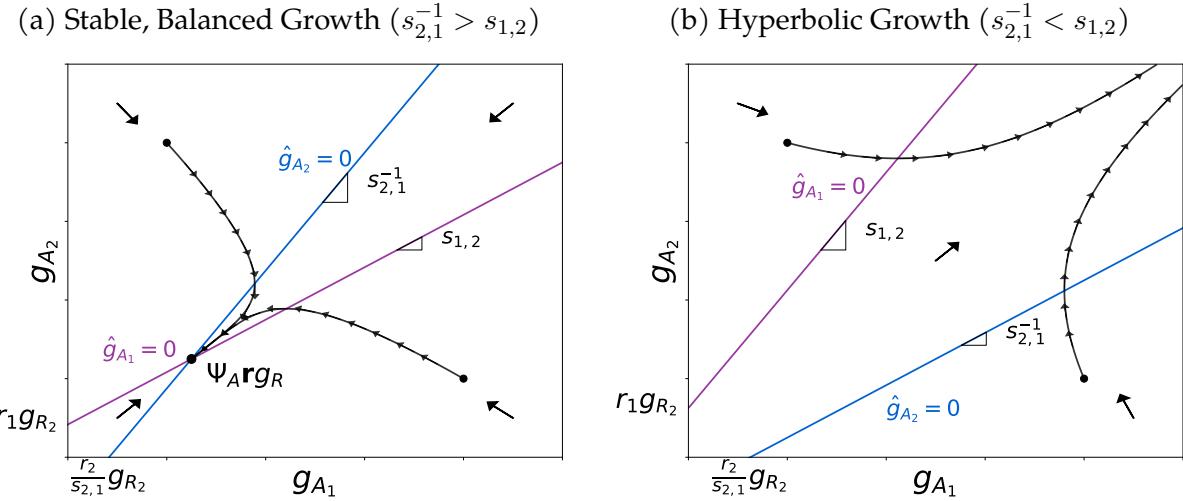
Note: Along colored lines the growth rate one of the technologies is constant (i.e., the growth-in-growth rate is zero). When these lines intersect the system exhibits stable balanced growth.

sector i, while β_i measures the strength of the sector-i "ideas-getting-harder-to-find" effect. Thus, explosive growth is more likely when either (1) spillovers are *large*, or (2) the ideas-getting-harder-to-find effect is *small*.

3.2 Economic and technological feedback loops

Now we introduce the innovation network of (17) into a broader economic environment. Specifically, instead of fixing the growth rate of research inputs as above, we endogenize this growth rate through a lab equipment model with exogenous population growth. We also allow technological progress to increase productivity in the final goods sector. In turn, this will lead to faster capital accumulation and ultimately lead to faster technological progress.

Specifically, consider the following system:

$$Y = \bar{A}K_Y^{1-\alpha}L_Y^{\alpha} \tag{22}$$

$$\dot{A}_i \propto (K_i^{1-\gamma_i}L_i^{\gamma_i})^{\lambda_i}A_i^{1-\beta_i}\prod_{j\in I\backslash i}A_j^{\phi_{i,j}} \tag{23}$$

$$\bar{A} = \prod_{i\in I}A_i^{\tau_i} \quad\text{where}\quad \sum_{i\in I}\tau_i = 1 \tag{24}$$

$$\dot{K} = aY - \delta K \tag{25}$$

$$K = K_Y + \sum_{i\in I}K_i \tag{26}$$

$$L = L_Y + \sum_{i\in I}L_i \tag{27}$$

where, a la Solow, we assume that the share of capital and labor allocated to each technology and output remain constant and a is a constant savings rate.

To write this in matrix form, define:

$$\mathbf{F}^Y := \frac{1}{\alpha}[\mathbf{r}(\mathbf{1}_N - \gamma)]'\tau \tag{28}$$

with $(\mathbf{1}_N - \gamma)$ and τ being the column vectors of capital contributions to research and the technology contributions to total factor productivity.

A balanced growth path in this environment can be characterized as follows.

Corollary 1 (Economic and Technological Feedback)**.** *The balanced growth path of technologies in the system described in equations* (22)*-*(27) *is given by*

$$\boxed{g_A^{BGP} = \mathbf{\Psi}_{A,Y}\mathbf{r}n} \tag{29}$$

where $\mathbf{\Psi}_{A,Y} = \left(\mathbf{I} - [\mathbf{F}^A + \mathbf{F}^Y]\right)^{-1}$ *and* $\mathbf{F}^A$ *and* $\mathbf{F}^Y$ *are the $I \times I$ technological and economic feedback matrices, with* $\mathbf{F}^A$ *and* $\mathbf{r}$ *defined above (equations* (20) *-*(21) *).*

From this balanced growth definition, we can immediately see that including economic, alongside technological, feedback loops necessarily brings the system closer to the hyperbolic tipping point; since $\mathbf{F}^Y$ is non-negative then $\mathbf{F}^Y + \mathbf{F}^A \geq \mathbf{F}^A$ entry-wise.

Note that entries in $\mathbf{F}^Y$ have a straightforward interpretation. The vector of terms, $\mathbf{r} \times (\mathbf{1}_N - \gamma)$, mediates the effect of output and hence capital on research inputs; on the

balanced growth path (with $g_Y^{\text{BGP}} = g_K^{\text{BGP}}$), from equation (23) we can arrive at

$$g_{A_i}^{\text{BGP}} = r_i(1 - \gamma)g_Y^{\text{BGP}} + \sum_{j \in I \setminus i} s_{i,j} g_{A_j}^{\text{BGP}} + r_i \gamma_i n.$$

Further, the vector of terms $\frac{1}{\alpha}\tau$ mediates the effect of technological progress on production; on the balanced growth path, from equation (22) we can arrive at

$$g_Y^{\text{BGP}} = \frac{1}{\alpha}\tau \cdot g_A^{\text{BGP}} + n.$$

Therefore, we can observe that entries in the feedback matrix exhibit complementarity: intensifying feedback in one direction of the output-technology loop amplifies the impact of strengthening feedback in the opposite direction.

In Figure 7 we illustrate in a two-sector example how feedback loops arise (separately) out of the output to technology vector (purple) and the technology to output vector (blue) from $\mathbf{F}^Y$, as well as the technology spillovers (green) from $\mathbf{F}^A$. Labor is also an input into each of these processes, but since it is non-accumulable and hence cannot participate in feedback loops we omit it from the figure.

Figure 7: Economic and Technological Feedback in a Two Technology Network
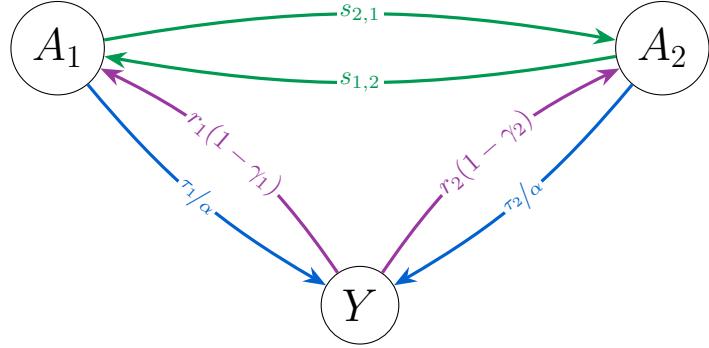
Note: Summarizing entries in $\boldsymbol{\Psi}_{A,Y}$, green lines represent the technology spillovers, $s_{i,j}$, that make up entries of $\mathbf{F}^A$; purple lines represent output-technology feedback mediated by the vector $\mathbf{r} \times (\mathbf{1}_N - \gamma)$ from $\mathbf{F}^Y$; blue lines represent the technology-output feedback mediated by the vector $\tau \times \alpha^{-1}$ from $\mathbf{F}^Y$.

The definition of the Leontief inverse, $\boldsymbol{\Psi}_{A,Y}$, provides a clear decomposition of the effects of technological and economic feedback loops to inform the balanced growth

path of this economy. Further, since this definition of the balanced growth path has an identical structure to that in the technology-only case, we can similarly apply Proposition 1:

Corollary 2. *Supposing population growth, n, is exogenous, finite, and positive, a balanced growth path exists for the system of equations in 22 - 27 if $\det\left(\mathbf{I} - [\mathbf{F}^A + \mathbf{F}^Y]\right) > 0$. Further, if $\det\left(\mathbf{I} - [\mathbf{F}^A + \mathbf{F}^Y]\right) < 0$ the system explodes in finite time.*

Just as with the balanced growth path under technological feedback loops alone, we can recover the standard balanced growth path from Jones (1995) if $\boldsymbol{\Psi}_{A,Y} = \mathbf{I}$. That is, (i) there are no technology spillovers, so $\mathbf{F}^A_{i,j} = 0$ for all i and j; and (ii) that no technology sector can simultaneously contribute to research (so $\gamma_i < 1$ for all i) and contribute to total factor productivity (so $\tau_i = 0$ for all i), so $\mathbf{F}^Y_{i,j} = 0$ for all i and j.

In this model we assume that savings rates are constant, but arbitrary. Of course, optimal savings consumption paths will respond endogenously to increases in the growth rate. However, such responses would only invalidate the above results if savings were to decline to zero in response to transformative growth induced by AI. Trammell and Korinek (2025) presents a simple argument using the Euler equation, illustrating that in a similar environment, given standard preferences, savings declines ultimately will not preclude explosive growth. In such a world, every unit of savings prior to the finite-time singularity can bring arbitrarily large returns at a future date, hence incentives to save become increasingly high.

4   AI and automation

The model above demonstrates the conditions for hyperbolic growth independent of automation. In this section we demonstrate that the networked-lab equipment model provides an intuitive framework to understand how AI and automation can accelerate growth: by strengthening pre-existing feedback loops or generating them in places they did not exist before. To develop this model, we suggest that AI will *replace labor* in some fraction of tasks in each research sector. Importantly, AI cognitive labor can accumulate in a way that human labor cannot; via both technological improvements and capital accumulation. We will see that by making assumptions on the process of automation we can calibrate the networked research model introduced above.

4.1 Automation and AI inputs

Here we adapt the automation framework from Zeira (1998), where *labor* outputs (be they from research or final production) are a Cobb-Douglas aggregation in outputs from tasks in the set, T. In particular, we assume that *effective* labor working in a sector i is given by

$$\hat{L}_i = p_i \prod_{q \in T} X_{i,q}^{\xi_q} \qquad \text{where} \sum_{q \in T} \xi_q = 1$$

where

$$X_{i,q} = \begin{cases} L_{i,q} & \text{if not automated} \\ C_{i,q} & \text{if automated} \end{cases}$$

where L_i is human labor and C_i is the human equivalent level of AI deployed on a task i and p_i is a productivity constant. I.e., we set C_i such that one unit of AI deployment to a task is equivalent to one unit of human labor. Assuming optimal allocation of AI equivalent labor and human labor (where AI and labor stocks are spread evenly across tasks), as the number of tasks approach infinity we have effective labor given by

$$\hat{L}_i(C_i, L_i) = \tilde{p}_i(f_i) C_i^{f_i} L_i^{1-f_i} \quad \text{where} \quad \tilde{p}_i(f_i) = \frac{p_i}{f_i^{f_i}(1-f_i)^{1-f_i}} \tag{30}$$

where $f_i \in [0, 1]$ is the fraction of tasks that are automated. Importantly, we treat f_i as an exogenous constant to evaluate the growth implications of a given level of task automation. Hence, our starting point is downstream of questions related to optimal automation of output or research tasks, which has received attention elsewhere (Jones, 2025, Acemoglu and Restrepo, 2018).

We separate AI capabilities, C_i, along two kinds of inputs:

- *Non-rivalrous*: progress in AI has been one of discovery of new, non-rivalrous, ideas. We assume that these non-rivalrous improvements result in some equivalent increase in AI equivalent labor. AI has emerged out of underlying progress in technologies such as compute hardware, algorithmic research or large-scale data collection.

- *Rivalrous*: the *use* of AI to complete tasks requires *inference compute* (compute to transform model inputs into productive work).

We emphasize that choice to model AI progress as arising from both *algorithmic* progress and from (inference) *compute scaling* is inspired by the crucial stylized fact that AI progress has been driven by both algorithmic advances and from increasing compute inputs (Ho et al., 2024).

In this spirit, we define our sectoral AI inputs as a function of a subset of technology sectors, $\bar{I} \subseteq I$; as well as a specific kind of capital, computing hardware, K^H.

Like in the rest of the model, we take the AI index to be a Cobb–Douglas function of a subset of underlying technologies, multiplied by the number of copies that can be run, a linear function of the amount of inference compute available in a given sector:

$$C_i = K_i^H \times \prod_{j \in \bar{I}} A_j^{\sigma_j} \tag{31}$$

where we do not necessarily assume that $\sum_{i \in \bar{I}} \sigma_j = 1$ since standard replication arguments for constant returns to scale don't apply to non-rivalrous technologies, A_i. We do assume that AI-equivalent labor scales linearly with inference compute.

4.2 Hyperbolic growth: Automated research

Here we demonstrate that we can recover the same structure on balanced growth as the general technological feedback equation (17) for a baseline semi-endogenous technological law of motion with automation. Hence, we can apply Proposition 1 to determine the conditions for hyperbolic growth, in an automated research environment.

We begin by assuming that the only inputs to research are *effective* labor, combining both human labor and AI effective labor according to equation (30). Therefore the technology laws of motion become

$$
\begin{aligned}
\dot{A}_i &= \nu_i \hat{L}_i^{\lambda_i} A_i^{1-\beta_i} \prod_{j \in I \neq i} A_j^{\phi_{i,j}} \\
&= \nu_i \tilde{p}_i(f_i) C_i^{f_i \lambda_i} L_i^{(1-f_i)\lambda_i} A_i^{1-\beta_i} \prod_{j \in I \neq i} A_j^{\phi_{i,j}} \\
&\propto \underbrace{A_i^{1-\beta_i}}_{\text{Dim. returns}} \times \underbrace{[K_i^{H f_i} L_i^{(1-f_i)}]^{\lambda_i}}_{\text{Riv. inputs}} \times \underbrace{\prod_{j \in I} A_j^{f_i \lambda_i \sigma_j}}_{\text{AI feedback}} \times \underbrace{\prod_{j \in I \neq i} A_j^{\phi_{i,j}}}_{\text{Direct spillovers}}
\end{aligned}
\tag{32}
$$

where $\sigma_j = 0$ for $j \notin \bar{I}$. As a baseline, we assume that direct technological spillovers are non-existent ($\phi = 0$). In this case, we recover sectoral spillovers through research

23

automation; a positive automation shock to a sector that feeds into AI progress accelerates research in all other sectors that have any automation. By matching terms we can calibrate our general network model with parameters we present the calibration in Table 1.

Table 1: AI-Parameterized R&D Network Model

General Model Input	AI-Parameterization	Description
β_i	$\beta_i - f_i\lambda_i\sigma_i$	Diminishing returns (offset)
$\phi_{i,j}$ (for $i \neq j$)	$f_i\lambda_i\sigma_j$	Automation spillovers
γ_i	f_i	Non-automated research task share
λ_i	λ_i	Parallelization penalty

Note: This table summarizes how one can re-parameterize the networked technology model (equation (17)). Blue terms (first three lines) represent additional terms arising from technological feedback loops. The final line has the same interpretation as in the general model.

From this calibration, we make several observations:

- Diminishing returns to research within a sector, β_i, are directly offset by the ability for AI to contribute to research in that sector. Introducing automated research is isomorphic to offsetting diminishing returns.

- For progress in j to spillover into progress in i, we require *both* automation in i ($f_i > 0$) and j to be relevant to AI progress ($\sigma_j > 0$). Therefore, a 'bilateral' feedback loop exists between i and j if both f and σ are greater than zero for both i and j. An 'indirect' feedback loop exists if there is a chain of technologies with positive f and σ such that progress in one sector eventually reinforces itself after accelerating technological progress in one sector, which subsequently accelerates in another sector and so on until the original sector is accelerated.

- If we were to deviate from our baseline assumption of no non-AI spillovers ($\phi_{i,j} > 0$), those spillovers would just be additive to automation spillovers: $\hat{\phi}_{i,j} = \phi_{i,j} + f_i\lambda_i\sigma_j$. In this case automation amplifies existing spillovers.

In summary, starting from a baseline of no sectoral spillovers and no contribution from lab-equipment, introducing research automation allows us to recover a networked, technological law of motion that conforms to a lab-equipment specification with spillovers. This means that after research automation, explosive conditions can emerge out of a model that precluded this possibility.

Substituting our recovered parameters into the balanced growth path described by equation (21) and denoting automation-adjusted parameters with hats, we have

$$g_A^{\text{BGP}} = \hat{\mathbf{\Psi}}_A \hat{\mathbf{r}} \hat{g}_R$$

where

$$
\begin{aligned}
\hat{g}_{R,i} &= f_i \times n + (1 - f_i) \times g_{H,i} \\
\hat{r}_i &:= \frac{\lambda_i}{\beta_i - f_i \lambda_i \sigma_i} \\
\hat{\mathbf{\Psi}}_A &= (\mathbf{I} - \hat{\mathbf{F}}_A)^{-1} \\
\hat{\mathbf{F}}_{i,j}^A &=
\begin{cases}
\hat{s}_{i,j} := \frac{f_i \lambda_i \sigma_j}{\beta_i - f_i \lambda_i \sigma_i} & \text{for } i \neq j \\
0 \text{ if } i = j
\end{cases}
\end{aligned}
$$

and black terms in $\hat{r}$ and $\hat{g}_R$ are those present in the baseline Jones (1995) model and while the blue terms in $\hat{r}$, $\hat{s}$ and $\hat{g}_R$ are those that enter through automated research channels.

We can apply Proposition 1 directly to this system. Yielding the result:

Corollary 3. *The automation-calibrated technology system (equation (32)) explodes in finite time iff*

$$\sum_{i \in I} r_i f_i \sigma_i > 1 \ .$$

Where this result comes directly from calculating $\det(\mathbf{I} - \hat{\mathbf{F}}^A)$ and rearranging.

4.3 Hyperbolic growth: Automated production and research

Now we extend the AI-induced feedback loops captured in equation (32) to a more realistic setting where AI can additionally automate some tasks in the production of goods, as well as allowing automated technological progress increase total factor productivity.

$$Y = \bar{A}K^{Y1-\alpha}\hat{L}_Y^{\alpha} \tag{33}$$

$$\bar{A} = \prod_{i\in I} A_i^{\tau_i} \quad \text{where} \quad \sum_{i\in I} \tau_i = 1 \tag{34}$$

$$\dot{A}_i \propto \hat{L}_i^{\lambda_i} A_i^{1-\beta_i} \tag{35}$$

$$\hat{L}_i \propto L_i^{1-f_i}C_i^{f_i} \quad \text{where} \quad C_i = K_i^H \times \prod_{j\in I} A_j^{\sigma_j} \tag{36}$$

$$\dot{K}^Y = a_Y Y - \delta_Y K^Y \tag{37}$$

$$\dot{K}^H = a_H Y - \delta_H K^H \tag{38}$$

$$K^H = K_Y^H + \sum_{i\in I} K_i^H \tag{39}$$

$$L = L_Y + \sum_{i\in I} L_i \tag{40}$$

where $a_H + a_Y \le 1$ and allocations of hardware and labour across research sectors is constant over time. Solving for the balanced growth path, we have

$$g_A^{\text{BGP}} = \hat{\boldsymbol{\Psi}}_{A,Y}\hat{\mathbf{r}}n \tag{41}$$

where

$$\hat{\boldsymbol{\Psi}}_{A,Y} = (\mathbf{I} - [\hat{\mathbf{F}}^A + \hat{\mathbf{F}}^Y])^{-1} \tag{42}$$

$$\hat{\mathbf{F}}^Y := \underbrace{[\hat{\mathbf{r}}\odot f]'}_{dg_A^{\text{BGP}}/dg_Y^{\text{BGP}} \text{ from } 35} \underbrace{\frac{1}{1-f_Y}[\frac{1}{\alpha}\tau + f_Y\sigma]}_{dg_Y^{\text{BGP}}/dg_A^{\text{BGP}} \text{ from } 33} \tag{43}$$

and $\hat{\mathbf{r}}$ and $\hat{\mathbf{F}}^A$ are defined as above.

In Table 2 we present how the core dynamics of the general R&D and production network model described by equations (22)-(27) can be calibrated to the automated economy model in equations (33) - (40) by adjusting relevant parameters. From this, we can see that introducing automation into the production side of the model (in addition to the technological side) is equivalent to: decreasing the labor share of production by a factor of $(1 - f_Y)$; and increasing the contributions of each technology to the production of final goods by $f_Y\sigma$.

Applying Proposition 1 to the balanced growth condition from equation (72), we

Table 2: AI-Parameterized R&D and Production Network Model

General Model Input	AI-Parameterization	Description
$1 - \alpha$	$1 - \alpha(1 - f_Y)$	Capital share of output
τ_i	$\tau_i + \sigma_i f_Y$	Technological contributions to output
β_i	$\beta_i - f_i \lambda_i \sigma_i$	Diminishing returns (offset)
$\phi_{i,j}$ (for $i \neq j$)	$f_i \lambda_i \sigma_j$	Automation spillovers
γ_i	f_i	Non-automated research task share
λ_i	λ_i	Parallelization penalty

Note: This table summarizes how one can re-parameterize the general model (equations (22)-(27)) to account changes in the parameter space due to automation. Purple terms (first two rows) represent additional terms arising from economic feedback loops while blue terms (middle three rows) represent additional terms arising from technological feedback loops. The final row has the same interpretation as in the general model.

have the following result:

Corollary 4. *With economic feedback loops, the automation-calibrated growth model (described in equations (33) - (40)) explodes in finite time iff*

$$f_Y + \sum_{i \in I} r_i f_i \left(\frac{\tau_i}{\alpha} + \sigma_i \right) > 1 \ . \tag{44}$$

Note, in Appendix B we derive the equivalent condition for the case of a fixed factor entering output, so output exhibits *diminishing* returns to capital and labor.

5 A calibrated application: Hyperbolic growth under AI-driven automation

Above we have introduced a general model of networked growth, where motivated these models with application to the case of automation of both final goods production and research. Importantly, the balanced and hyperbolic growth conditions in section

[4](#) are stated in terms of parameters that can be directly calibrated, based on historical evidence on diminishing returns in software and hardware research together with the labor share of output.

5.1 Simple integrated AI-economy

Here we integrate a simple model of AI progress within an economic environment. As in previous sections, we assume that AI contributes to cognitive labor, replacing human labor in some fraction of tasks in each sector. The central force continues to be that AI progress stems from both better algorithms and better computing hardware. Improved computing hardware allows us to run more computations for the same amount of capital investment, while improved algorithms make AI more capable of completing relevant tasks.

We deliberately omit a number of additional components of the training process here to emphasize key feedback loops. For example, our model exclusively focuses on *inference* without modeling the relationship between training investment and inference capabilities. Erdil et al. (2025) considers a much richer AI-economic model, though such a framework makes it impossible to cleanly isolate the key feedback loops of interest.

We present the equations of the model before describing in words:

$$Y = A\hat{L}_Y^\alpha K^{Y^{1-\alpha}} \tag{45}$$

$$\hat{L}_i \propto L_i^{1-f_i} C_i^{f_i} \quad \text{where} \quad C_i = K_i^H S \tag{46}$$

$$\dot{S} \propto \hat{L}_{\dot{S}}^{\lambda_S} S^{1-\beta_S} \tag{47}$$

$$\dot{H} \propto \hat{L}_{\dot{H}}^{\lambda_H} H^{1-\beta_H} \tag{48}$$

$$\dot{A} \propto \hat{L}_{\dot{A}}^{\lambda_A} A^{1-\beta_A} \tag{49}$$

$$\dot{K}^H = a_H HY - \delta_H K^H \tag{50}$$

$$\dot{K}^Y = a_Y Y - \delta_Y K^Y \tag{51}$$

$$K^H = K_Y^H + K_{\dot{H}}^H + K_{\dot{S}}^H + K_{\dot{A}}^H \tag{52}$$

$$L = L_Y + L_{\dot{H}} + L_{\dot{S}} + L_{\dot{A}} \tag{53}$$

In this environment, we have three independent technological processes (general TFP A, software S, and hardware quality H) which amplify feedback loops. Feedback from output Y to effective labor $\hat{L}$ through accumulation of compute K^H supports the infer-

ence of AI models. This channel from output to compute is amplified through hardware quality progress; we assume hardware progress means more inference compute can be purchased for the same amount price over time in (50). The capacity of effective labor to contribute to both research and production is amplified through software progress S, which makes compute more effective at completing economic tasks. Here we define software in terms of productivity units of human labor and that doubling the software 'level' means that software can produce double the output on a specific task given the same amount of inference compute. The role of capital is unchanged from standard models, and labor grows exogenously.

Just as in Section 4.3, we can find the conditions for balanced and explosive growth by recognizing from equations (50) and (51) that a balanced growth path requires $g_{KY} = g_{KH} - g_H = g_Y$. Then we can use the equations for output and effective labor to solve for the balanced growth path of technologies as a function of fundamentals of the model. This is the same balanced growth condition for AI capabilities as in the general model from Section 4.3. Further, note that we only have one technology – TFP – feeding *directly* into increasing output productivity.

Therefore, we can see that the system described in equations (47)-(53) is a three technology version of the general model from equations (33)-(40), where: technological contributions to output are dictated by $\tau_z = \tau_h = 0$ and $\tau_A = 1$; and technological contributions to AI progress are dictated by $\sigma_z = \sigma_h = 1$ and $\sigma_A = 0$.[9]

Given that the AI-economic model developed here is a specific case of the general model, we can simply calibrate Corollary 4 to derive the hyperbolic growth condition:

Corollary 5. *The AI-economic model (described in equations (45)-(53)) explodes in finite time iff*

$$f_Y + f_S r_S + f_H r_H + f_A \frac{r_A}{\alpha} > 1 \tag{54}$$

5.2 Calibration

We now turn to calibrating the relevant terms from equation (54). In Table 3 we report estimates of research productivity in software, hardware and aggregate TFP. Both software and (in particular) hardware research are significantly more productive than

---

[9]The only difference between this calibrated model and the general model from Section 4.3 is that one of the technologies – hardware quality – scales compute accumulation rather than AI capabilities directly. Ultimately this does not affect the balanced growth path calculation.

29

aggregate TFP.

Term	Parameter	Estimate	Source
Labor share	α	0.6	
Returns to research (software)	r_S	~ 1	Erdil, Besiroglu and Ho (2024)
Returns to research (hardware)	r_H	5	Bloom et al. (2020)
Returns to research (TFP)	r_A	0.32	Bloom et al. (2020)

We postpone discussion of the limitations of this calibration to Section 5.3.

Table 4 presents the calibrated explosion conditions, i.e. from calibrating (54) using table 3. The first row presents the condition if the software sector is the only sector to be (partially) automated; the second, if only hardware; and so on, with different variations.

A key takeaway of the exercise is that automating hardware research – increasing f_H – has the highest impact of automation across any sector, since the returns to research are so high in that sector, $r_H = 5$. Automating one hardware research task offers about the same increase in the distance to the threshold as five tasks in software or final goods production and the same as ten tasks in the general TFP sector.

Further, we can see that introducing a 10% fixed factor so that output is diminishing returns to scale in capital and effective labor, there is only a mild effect on the threshold since this effect does not pass through the hardware or software automation channels.[10] Appendix B provides the analytical balanced and hyperbolic growth conditions with a fixed factor.

Table 5 first solves for how much automation is required to achieve a doubling of output growth on the balanced growth path, assuming an equal level of automation in each sector under consideration, if feasible. Column 2, meanwhile, solves for the level of automation required to achieve the hyperbolic growth threshold, under the same assumption of equal automation across sectors.

We can see that in the full model, the automation required to double the balanced growth path is about three quarters of the automation necessary to achieve fully ex-

---

[10]Fixed factors can be introduced in the above model by redefining output shares on labor and capital according to $\alpha_L = 0.9 \times \alpha$ and $\alpha_K = 0.9 \times (1 - \alpha)$.

Table 4: Applying historical estimates of research productivity to singularity conditions

Sectors with Automation	Calibrated Explosion Condition
S	$\sim 1f_S > 1$
H	$5f_h > 1$
S, H	$\sim 1f_S + 5f_h > 1$
H, Y	$5f_H + f_Y > 1$
S, Y	$\sim 1f_H + f_Y > 1$
S, H, Y	$\sim 1f_S + 5f_h + f_Y > 1$
S, H, A, Y	$\sim 1f_S + 5f_h + 0.53f_A + f_Y > 1$
S, H, A, Y (10% fixed factor)	$\sim 1f_z + 5f_h + 0.45f_A + 0.86f_Y > 1$

Note: Estimates of historical returns to research effort are taken from Bloom et al. (2020) for total factor productivity and hardware, and from Erdil, Besiroglu and Ho (2024) for software. Here, f terms can be calibrated to the fraction of tasks that can be automated in research sectors; software (S), hardware (H), and general technology A.

plosive growth. That is, the balanced growth path is initially very slow to respond to changes in automation, and then changes incredibly quickly. This is the result of the multiplication of spillovers through the Leontief inverse.

These quantitative results underscore two points. First is the importance of amplification of feedback loops. Under the simple model developed above, higher levels of automation – dialing up feedback loops – accumulates into rapid changes in the growth path quickly. Second is the emphasis that these are changes in the asymptotic balanced growth path. In reality it may take some time to converge to such a path and further, even under a hyperbolic growth path it may still take significant time to actually see radical changes in growth rates; hyperbolic growth does not necessarily imply transformative growth in the short run. One can see this intuitively by recognizing that from the view point of the entire human history the current 2% annual growth *is* part of a long-run super-exponential trend.[11]

---

[11] Hanson (2000) makes this point through a theory-free projection of growth by fitting historical growth rates to a series of increasing exponential functions until 1998 which predicted growth rates

Table 5: Automation for Hyperbolic vs Balanced Growth

Automated Factors	Automation Threshold	
	$2 \times g_Y^{\text{BGP}}$	Hyperbolic
S	–	$\sim 100\%$
H	–	20%
S, H	–	17%
H, Y	12%	17%
S, Y	16%	50%
S, H, Y	11%	14%
S, H, A, Y	8%	13%
S, H, A, Y (10% fixed factor)	9%	14%

Note: Here we assume that $f = f_S = f_H = f_A = f_Y$. In the middle column we solve for the f such that the balanced growth path doubles relative to no automation. In the right column we solve for the f such that the conditions from Table 4 are satisfied. The first three balanced growth rows are empty since these are technology-only automation scenarios.

5.3   Calibration limitations

Here we offer an illustrative calibration of the explosion threshold based on historical estimates of research productivity in different sectors. However, there are a number of reasons why these estimates might be inappropriate for this model.

First, is that returns to research effort in software from Erdil, Besiroglu and Ho (2024) have been estimated in software domains other than frontier AI research. Specifically, these estimates come from chess engines and computer vision. Estimating r_S directly from the rate of progress at frontier AI labs is difficult not just because of challenges associated with estimating research inputs, but also because r_S is estimated as returns to research along the balanced growth path. Given the boom in progress in AI it is impossible to tell if we are anywhere close to balanced growth software research. Ho and Whitfill (2025) make some attempt to make this calculation directly, finding r_S

___
larger than 20% by 2040.

in the range of 1.2 - 1.8. Though, given limitations in this approach, we rely on relatively conservative estimates from Erdil, Besiroglu and Ho (2024).

Second, we assume that these parameters are foundational to the knowledge discovery process, rather than the *human* knowledge discovery process. For example, that diminishing returns to research (β) or parallelization of research (λ) are independent of whether it is humans of AI completing that scientific research. Trammell and Korinek (2025) suggest reasons why both of these parameters may be different under AI R&D. For example, as identified by Ekerdt and Wu (2025), increasing the researcher share of population may result in declines in returns to research effort as the average quality of researcher declines; in the case of AI we might expect constant effective researcher quality, increasing r for AI researchers relative to human researchers.

Third, we take estimates from Bloom et al. (2020) and Erdil, Besiroglu and Ho (2024) as given. Specifically, the estimate of r_A from Bloom et al. (2020) is an aggregate estimate of returns to research across the whole economy. However, we should expect that the research that has contributed to hardware and software progress have in fact contributed to aggregate TFP. Since we separate out software and hardware progress from TFP we should also adjust r_A to be the returns to research in TFP, net of software and hardware sectors. We do not have a good estimate of the share of aggregate technological progress that has come from AI software and hardware research (τ_S and τ_H terms from above) hence we just assume these are small and calibrate r_A as the aggregate economy estimate.

6   Discussion

The above analysis employs a simple model to derive a critical threshold for the automation of tasks: once this threshold is crossed, intelligence and output are projected to grow super-exponentially, ultimately reaching infinity within finite time. Crucially, the possibility of a singularity hinges on two assumptions: constant research productivity up to the singularity and correctly specified research production functions, which may be satisfied soon, even though they may fail to hold in the long run.

Regarding constant research productivity, our findings indicate that automating hardware quality research is a significant contributor to meeting the explosion condition, largely due to the extreme pace of Moore's law over the past half-century. However, physical constraints on increasing transistor density in chips that support AI train-

ing and inference may limit this progress. Current state-of-the-art chips already feature transistors as small as 2 nanometers, which may approach the boundaries of physical feasibility. Such challenges have fostered skepticism regarding the continued validity of Moore's law, which has historically predicted consistent advancements in transistor density (Leiserson et al., 2020).

Regarding the specification of the production function, we have assumed that the production factors – human labor and AI cognitive labor – are substitutes. In the short run, this assumption may hold; for instance, AI is already capable of performing coding tasks semi-autonomously, replacing human programmers. However, as automation advances, human and AI cognitive labor may appear complementary over some time periods. For example, while AI may excel in chip architecture, manufacturing, and testing, it might remain unable to conduct the physical experiments necessary for chip design before sufficiently advanced robots are developed, thereby relying on human researchers. Such considerations receive substantial consideration in Jones (2025). In such cases, the pace of hardware improvement could be constrained by the availability of human experimenters. By alleviating reliance on this scarce resource, these efforts could delay or mitigate the stalling of progress caused by bottlenecks. While such adaptive responses may not eliminate bottlenecks entirely, they could extend the period of rapid advancements before fundamental constraints take hold.

Although limitations are likely to emerge in the long run, we argue that assuming (roughly) constant research productivity and substitutable labor provides a reasonable approximation of the production process in the short term. Thus, a more realistic interpretation of the explosion conditions derived above is that they signify a threshold for temporary super-exponential growth take-off in economic output and AI capabilities. Crucially, this super-exponential growth is driven not by resource reallocation within the economy but by changes in the fundamental drivers of economic and technological progress.

7   Conclusion

This paper develops a framework for understanding how advances in artificial intelligence could fundamentally transform economic growth dynamics. By modeling the interconnected roles of hardware, software, and general technological progress, we show that automation of research and development activities could generate powerful

feedback effects leading to rapid growth acceleration. Our calibration using historical estimates of research productivity suggests that these effects could be substantial, particularly through the automation of semiconductor research.

These findings have important implications for economic policymakers. First, they suggest that standard growth models may need significant revision to account for the potential of recursive technological progress. Traditional frameworks that treat technological advancement as an exogenous or smoothly evolving process may not capture the possibility of sudden acceleration in growth rates driven by AI automation of research activities.

Second, our results highlight the strategic importance of semiconductor research and development. The high historical productivity of hardware research, combined with growing automation capabilities in chip design, suggests that this sector could play a crucial role in determining the pace of overall technological progress. This raises important questions about both the market concentration and geographic concentration of semiconductor research and production capabilities (Korinek and Vipra, 2025).

Third, our analysis suggests that monitoring automation levels in research and development activities may be as important as tracking traditional macroeconomic indicators. The extent of automation in key research sectors could serve as an early warning system for potential growth acceleration.

Several promising directions for future research emerge from our analysis. Researchers could develop systematic metrics to measure automation levels across different research domains. Empirical work might test for the presence and strength of feedback loops between AI advances and technological progress, particularly in semiconductor and software research. Future studies could investigate the degree of complementarity between human and AI researchers in different activities. Analysis of how existing concentrations of AI and semiconductor research capabilities might influence regional growth patterns would be valuable. Finally, work is needed to better characterize the physical, computational, and economic constraints that might prevent growth acceleration.

Understanding these dynamics is crucial for economists and economic policymakers. If AI progress can indeed generate self-reinforcing technological acceleration, this has profound implications for economic planning, research funding priorities, and international economic coordination. While our analysis suggests this scenario is plausible given historical patterns of research productivity, much work remains to be done in understanding how these dynamics might unfold in practice.

References

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.

Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones. 2019. *Artificial Intelligence and Economic Growth.* University of Chicago Press.

Altman, Sam. 2025. "The Gentle Singularity." *Blog post*.

Besiroglu, Tamay, Nicholas Emery-Xu, and Neil Thompson. 2024. "Economic impacts of AI-augmented R&D." *Research Policy*, 53(7): 105037.

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

Cottier, Ben, Robi Rahman, Loredana Fattorini, Nestor Maslej, Tamay Besiroglu, and David Owen. 2024. "The rising costs of training frontier AI models." *arXiv preprint arXiv:2405.21015*.

Ekerdt, Lorenz, and Kai-Jie Wu. 2025. "Self-Selection and the Diminishing Returns of Research."

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."

Erdil, Ege, Tamay Besiroglu, and Anson Ho. 2024. "Estimating Idea Production: A Methodological Survey." *arXiv preprint arXiv:2405.10494*.

Goldie, Anna, and Azalia Mirhoseini. 2024. "How AlphaChip Transformed Computer Chip Design." *Google DeepMind Blog*.

Good, Irving John. 1966. "Speculations Concerning the First Ultraintelligent Machine." In . Vol. 6 of *Advances in Computers*, , ed. Franz L. Alt and Morris Rubinoff, 31–88. Elsevier.

Hanson, Robin. 2000. "Long-term growth as a sequence of exponential modes." *Working manuscript*.

Ho, Anson, and Parker Whitfill. 2025. "The software intelligence explosion debate needs experiments." *Epoch AI Substack (Gradient Updates)*.

Ho, Anson, Tamay Besiroglu, Ege Erdil, David Owen, Robi Rahman, Zifan C Guo, David Atkinson, Neil Thompson, and Jaime Sevilla. 2024. "Algorithmic progress in language models." *Advances in Neural Information Processing Systems*, 37: 58245–58283.

Jones, Ben. 2025. "Artificial Intelligence and Research and Development." National Bureau of Economic Research.

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

Jones, Charles I., and Christopher Tonetti. 2025. "Past Automation and Future A.I.: How Weak Links Tame the Growth Explosion." Working paper, Stanford GSB and NBER.

Korinek, Anton, and Jai Vipra. 2025. "Concentrating Intelligence: Scaling and Market Structure in Artificial Intelligence." *Economic Policy*. Forthcoming.

Leiserson, Charles E, Neil C Thompson, Joel S Emer, Bradley C Kuszmaul, Butler W Lampson, Daniel Sanchez, and Tao B Schardl. 2020. "There's plenty of room at the Top: What will drive computer performance after Moore's law?" *Science*, 368(6495): eaam9744.

Liu, Ernest, and Song Ma. 2024. "Innovation networks and r&d allocation." National Bureau of Economic Research.

Mirhoseini, Azalia, Anna Goldie, Mustafa Yazgan, Joe Wenjie Jiang, Ebrahim Songhori, Shen Wang, Young-Joon Lee, Eric Johnson, Omkar Pathak, Azade Nazi, et al. 2021. "A graph placement methodology for fast chip design." *Nature*, 594(7862): 207–212.

Nesov, Vladimir. 2025. "Slowdown After 2028: Compute, RLVR Uncertainty, MoE Data Wall." *LessWrong*.

Ngai, L Rachel, and Roberto M Samaniego. 2011. "Accounting for research and productivity growth across industries." *Review of Economic Dynamics*, 14(3): 475–495.

OpenAI. 2025. "Sam, Jakub, and Wojciech on the future of OpenAI with audience Q&A." *YouTube livestream*, Featuring Sam Altman (CEO), Jakub Pachocki (Chief Scientist), and Wojciech Zaremba.

Trammell, Philip, and Anton Korinek. 2025. "Economic growth under transformative AI." National Bureau of Economic Research.

Whitfill, Parker, Ben Snodin, and Joel Becker. 2025. "Forecasting AI Time Horizon Under Compute Slowdowns." *arXiv preprint arXiv:2511.19492*.

Yudkowsky, Eliezer. 2013. "Intelligence explosion microeconomics." *Machine Intelligence Research Institute*, 23: 2015.

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

A   Deriving conditions for hyperbolic growth

Proof of Proposition 1. First we divide equation (17) by A_i to get technology growth rates and then take the logs and time derivative to get the rate of change in growth rates given by the vector

$$\dot{g}_A = diag(g_A)[(\mathbf{S} - \mathbf{B})g_A + diag(\ell)g_E]$$

where $\mathbf{S}_{i,j} = p_{i,j}$ for $i \neq j$ and zero for diagonal elements, and $\mathbf{B} = diag(b_1, b_2, \ldots, b_N)$. We define the 'exponent matrix', $\mathbf{\Omega} := \mathbf{S} - \mathbf{B}$. Therefore, We can relate the (balanced growth path) spillover matrix, $\mathbf{F}$ and the exponent matrix $\mathbf{\Omega}$ according to $\mathbf{\Omega} = \mathbf{B}(\mathbf{F} - \mathbf{I})$. Further, let $u \gg 0$ be the Perron–Frobenius right eigenvector of $\mathbf{F}$, $\mathbf{F}u = \rho(\mathbf{F})u$, and set $\mu := \rho(\mathbf{F}) - 1$. Then

$$(\mathbf{F} - \mathbf{I})u = \mu u, \qquad \mathbf{\Omega}u = \mathbf{B}(\mathbf{F} - \mathbf{I})u = \mu\,\mathbf{B}u \quad (\gg 0).$$

Further, we define the scalars

$$\bar{h}(t) := \frac{g_{A_{\bar{i}}}(t)}{u_{\bar{i}}} \quad (> 0), \quad \text{where} \quad \bar{i} = \arg\max_i \frac{g_{A,i}(t)}{u_i}$$

$$\underline{h}(t) := \frac{g_{A_{\underline{i}}}(t)}{u_{\underline{i}}} \quad (> 0), \quad \text{where} \quad \underline{i} = \arg\min_i \frac{g_{A,i}(t)}{u_i}$$

Proving $\rho(\mathbf{F}) > 1 \implies$ hyperbolic growth. Suppose $\rho(\mathbf{F}) > 1$ so $\mu > 0$. Take the rate of change in the growth rate of $A_{\underline{i}}$ at a specific time t.

$$\dot{g}_{A,\underline{i}}(t) = g_{A,\underline{i}}(t)(\mathbf{\Omega}g_A + diag(\ell)g_E)_{\underline{i}} \tag{55}$$

$$\geq g_{A,\underline{i}}(t)(\mathbf{\Omega}g_A)_{\underline{i}} \tag{56}$$

$$= g_{A,\underline{i}}(t)\left(\sum_{j \neq \underline{i}} \mathbf{\Omega}_{j,\underline{i}}g_{A,j} + \mathbf{\Omega}_{\underline{i},\underline{i}}g_{A,\underline{i}}(t)\right) \tag{57}$$

$$\geq g_{A,\bar{i}}(t)\left(h(t)\sum_{j \neq \underline{i}} \mathbf{\Omega}_{j,\underline{i}}u_j + \mathbf{\Omega}_{\underline{i},\underline{i}}g_{A,\underline{i}}(t)\right) \tag{58}$$

$$D = \underline{h}(t)^2 u_{\underline{i}}(\mathbf{\Omega}u)_{\underline{i}} \tag{59}$$

where 58 comes from the fact that by definition $g_{A,j}(t) \geq h(t)u_j$ as well as that $\Omega_{j,i} \geq 0$ and 59 comes from the fact that this holds with equality when $j = \underline{i}$. Next, since $\dot{\underline{h}}(t) =$

$\dot{g}_{A,\underline{i}}(t)/u_{\underline{i}}$ then we have

$$\dot{\underline{h}}(t) \;\geq\; \frac{h(t)^2}{u_{\bar{i}}}(\boldsymbol{\Omega}u)_{\bar{i}} \tag{60}$$

$$=\; \frac{h(t)^2}{u_{\bar{i}}}(\mu\mathbf{B}u)_{\bar{i}} \tag{61}$$

$$=\; \underline{h}(t)^2\mu b_{\underline{i}}u_{\underline{i}} \tag{62}$$

which implies $\underline{h}(t)$ grows hyperbolically, which in turn implies that g_{A_i} grows hyperbolically for all i by definition of $h(t)$, which implies A_i grows hyperbolically $\forall i$.

Next, proving the $\rho(F) \leq 1 \implies$ no hyperbolic growth. Assuming $\rho(\mathbf{F}) \leq 1$ gives $\mu \leq 0$. Following a similar procedure as above, we have

$$\dot{g}_{A,\bar{i}} \;=\; g_{A,\bar{i}}(t)(\boldsymbol{\Omega}g_A(t) + diag(\ell)g_E)_{\bar{i}} \tag{63}$$

$$=\; g_{A,\bar{i}}(t)(\boldsymbol{\Omega}g_A)_{\bar{i}} + g_{A,\bar{i}}(t)\ell_i g_E \tag{64}$$

$$=\; g_{A,\bar{i}}(t)\Big(\sum_{j\neq\bar{i}}\boldsymbol{\Omega}_{\bar{i},j}g_{A,j} + \boldsymbol{\Omega}_{\bar{i},\bar{i}}g_{A,\bar{i}}\Big) + g_{A,\bar{i}}(t)\ell_{\bar{i}}g_E \tag{65}$$

$$\leq\; g_{A,\bar{i}}(t)\Big(\sum_{j\neq\bar{i}}\boldsymbol{\Omega}_{\bar{i},j}\bar{h}(t)u_j + \boldsymbol{\Omega}_{\bar{i},\bar{i}}g_{A,\bar{i}}\Big) + g_{A,\bar{i}}(t)\ell_{\bar{i}}g_E \tag{66}$$

$$=\; g_{A,\bar{i}}(t)\bar{h}(t)\Big(\sum_{j\neq\bar{i}}\boldsymbol{\Omega}_{\bar{i},j}u_j + \boldsymbol{\Omega}_{\bar{i},\bar{i}}u_{\bar{i}}\Big) + \bar{h}(t)u_{\bar{i}}\lambda_i g_E \tag{67}$$

$$=\; \bar{h}(t)^2 u_{\bar{i}}(\boldsymbol{\Omega}u)_{\bar{i}} + \bar{h}(t)u_{\bar{i}}\ell_{\bar{i}}g_E \tag{68}$$

$$=\; \bar{h}(t)^2 u_{\bar{i}}\mu(\mathbf{B}u)_{\bar{i}} + \bar{h}(t)u_{\bar{i}}\ell_{\bar{i}}g_E \tag{69}$$

and since $\dot{\bar{h}}(t)u_{\bar{i}} = \dot{g}_{A,\bar{i}}$ then we can upper bound by the logistic differential equation

$$\dot{\bar{h}}(t) \leq \bar{h}(t)^2\mu\beta_{\bar{i}}u_{\bar{i}} + \bar{h}(t)u_{\bar{i}}\ell_{\bar{i}}g_R \tag{70}$$

and since $\mu < 0$ the quadratic part of the expression dominates as $\bar{h}$ grows so $\dot{\bar{h}}(t)$ remains finite for all t. Further in the case of $\mu = 0$ the inequality reduces to $\dot{\bar{h}}(t) \leq \bar{h}(t)\ell_{\bar{i}}g_E$, yielding at most exponential growth in $\bar{h}$. Finally, we know that no explosive growth in $\bar{h}$ implies there is no explosive growth in $g_{A,i}$ for all $i \in I$ (nor in A_i).

Finally we prove fully endogenous balanced growth with $\rho(\mathbf{F}) = 1$ and $g_E = 0$. In this case, we have the the motion of technology growth balanced growth path $g_A^{\text{BGP}} = \mathbf{F}g_A^{\text{BGP}}$. From above we have that $\mathbf{F}u = \rho(\mathbf{F})u$ and when $\rho(\mathbf{F}) = 1$ then $u = \mathbf{F}u$. Finally, if $u = \mathbf{F}u$ and $g_A^{\text{BGP}} = \mathbf{F}g_A^{\text{BGP}}$ then we require $g_A^{\text{BGP}} \propto u$. $\qquad\square$

B Balanced and hyperbolic growth with a fixed factor

Here we take the growth model from Section 4.3, but allow for the inclusion of a fixed factor, M, into output. Where output is constant returns to scale in capital, labor and the fixed factor

$$Y \;=\; \bar{A}K_Y^{\alpha_K}\hat{L}_Y^{\alpha_L}M^{1-\alpha_K-\alpha_L}. \tag{71}$$

We maintain equations (35) - (40) to describe law of motion of technology, hardware, capital and resource constraints.

Solving for the balanced growth path under the assumption that output is constant returns to scale in capital and labor we have

$$g_A^{\text{BGP}} \;=\; \hat{\boldsymbol{\Psi}}_{A,Y} \times \left(\hat{\mathbf{r}} \times \frac{\mathbf{1}_n - f - \alpha_K(\mathbf{1}_n - f) - \alpha_L(\mathbf{1}_n \times f_Y - f)}{1 - \alpha_K - \alpha_L f_Y}\right) \times n \tag{72}$$

where

$$\hat{\boldsymbol{\Psi}}_{A,Y} \;=\; (\mathbf{I} - [\hat{\mathbf{F}}^A + \hat{\mathbf{F}}^Y])^{-1} \tag{73}$$

$$\hat{\mathbf{F}}^Y \;:=\; \underbrace{[\hat{\mathbf{r}} \odot f]'}_{dg_A^{\text{BGP}}/dg_Y^{\text{BGP}} \text{ from tech L.O.M}} \times \underbrace{\frac{\alpha_L}{1 - \alpha_K - \alpha_L f_Y} \times [\frac{1}{\alpha_L}\tau + f_Y\sigma]}_{dg_Y^{\text{BGP}}/dg_A^{\text{BGP}} \text{ from } 71} \tag{74}$$

and $\hat{\mathbf{r}}$ and $\hat{\mathbf{F}}^A$ are defined as in Section 4.2.

Then, deriving the explosive growth condition, we get

Corollary 6. *The automation-calibrated growth model with a fixed factor (described in equations* (71) *and* (35) - (40)) *explodes in finite time iff*

$$\frac{\alpha_L}{1 - \alpha_K}f_Y + \sum_{i \in I} f_i r_i \left(\frac{\tau_i}{1 - \alpha_K} + \sigma_i\right) > 1 \tag{75}$$