When Does Automating AI Research Produce Explosive Growth?

Feedback Loops in Innovation Networks

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

May 2026

Abstract

AI labs are increasingly using AI itself to accelerate AI research, creating a feedback loop that could lead to an intelligence explosion. We develop a general semi-endogenous growth model with an innovation network, where research and automation in one sector increase the productivity of research in other sectors, and derive a clean analytical condition under which growth becomes superexponential ("explosive"). We find that automating research can offset diminishing returns to ideas by activating two reinforcing channels: a technological feedback loop across research sectors, and an economic feedback loop in which higher output finances further research. Growth becomes explosive if the combined strength of technological and economic feedback loops overcomes diminishing returns. In a simple simulation calibrated to trends in AI progress, fully automating software research and modest (5%) automation in other sectors generates a singularity within six years. Bottlenecks do not overturn the result if task automation advances sufficiently fast.

Keywords: AI, economic growth, automation, innovation networks, feedback loops, intelligence explosion, singularity

*Forethought: tom@forethought.org
†University of Virginia: basilh@virginia.edu
‡Columbia University: twh2125@columbia.edu
§University of Virginia and EconTAI, Brookings, NBER, and CEPR: econtai@virginia.edu
We thank Dan Carey, Philip Trammell, Cheryl Wu, Parker Whitfill and participants at the Oxford Global Priorities Workshop and a seminar at Chicago Booth for helpful comments and discussion.

1 Introduction

> "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

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 (Yudkowsky, 2013). OpenAI bluntly declares its goal of developing such technology "by March of 2028" (OpenAI, 2025).

Whether such recursive self-improvement actually produces explosive economic growth depends on two economic forces:

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 $\beta > 0$. 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). Intuitively, even if AI were 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. We develop a framework for analyzing how AI-driven automation interacts with both forces, and identify the conditions under which feedback loops generated by automation tip the economy into explosive growth. We begin with a benchmark model without bottlenecks to isolate whether automation can overturn diminishing returns in research; we then show how fast automation must progress for this result to survive when tasks are gross complements.

We begin by developing a general semi-endogenous growth model with multiple different technologies that may impose spillovers on each other across a network. This structure may be of independent interest to the growth literature but is particularly relevant for our analysis of the AI R&D process. We allow for automation in the ideas production function (Aghion, Jones and Jones, 2019) and examine how automation in one research sector spills through the innovation network and interacts with economic growth. We structure our argument by starting with a framework with fixed levels of automation and without bottlenecks, to focus on the role of diminishing returns; and then generalize to study bottlenecks when automation increases over time.

The model identifies two distinct channels through which automation generates explosive dynamics, and these channels mutually reinforce each other. The first is *technological* feedback loops across the innovation network. In the canonical semi-endogenous model, there is one research sector; our model instead 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 the design of 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. The second channel is an *economic* feedback loop, in which higher output generates more resources that can be deployed to drive further economic growth. The classic example is capital accumulation: higher output leads to more investment, which produces yet more output. In our setting, this channel is particularly important: AI R&D investment is expensive and growing tenfold every two years (Cottier et al., 2024; Whitfill, Snodin and Becker, 2025; Nesov, 2025), and higher output allows for continued such investment.

Analytical insights. The general model produces a simple analytical condition under which technological and economic feedback loops give rise to balanced or explosive growth, along with several broad, interpretable insights. First, technological feedback loops – spillovers across research sectors – can offset diminishing returns in the production of ideas. They reduce the degree to which ideas get harder to find, so that if such spillovers are improperly ignored, a system may be estimated to be non-explosive when in reality spillovers from other sectors tip it into explosive dynamics.

Second, automation creates economic feedback loops, which similarly counteract diminishing returns. Automation means that a task previously performed by human labor is instead performed by machines, i.e. capital. Human labor does not generate an economic feedback loop in the modern era: higher GDP does not result in a higher population. Machines, on the other hand, do: higher GDP leads to the construction of more machines. Replacing human labor with capital therefore counteracts the diminishing returns in the production of ideas by creating a new feedback loop where one did not exist previously.

Third, technological and economic feedback loops amplify each other. AI research results in higher output, which helps fund further investment in AI research, which in turn drives further technological progress.

Application: AI automation with software & hardware feedback loops. We apply the general framework to study an economy with AI-driven automation, 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 delivers closed-form results. Figure 1 provides an overview of the structure.

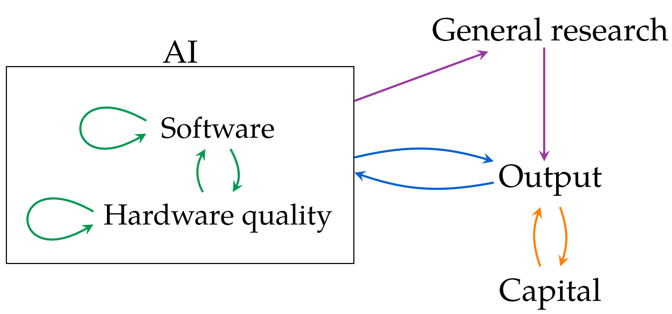

Figure 1: Model of AI automation with software & hardware feedback loops

Our starting assumption is that AI can automate a 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 or by AI. Finally, goods output likewise may be produced by a combination of humans and 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]

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

Here, the subscripts denote sectors: final output Y, software research S, hardware research H, and general research A. The term $f_i \in [0, 1]$ is the fraction of tasks in each sector that have been automated by AI. The term $\beta_i > 0$ is the degree of diminishing returns in each research sector. $1 - \alpha$ is the labor share in production of output.

The explosion condition (1) implies that growth is hyperbolic if the sum of the strengths of feedback loops exceeds unity. In particular, the relevant feedback loops are: (i) the pure economic feedback loop (which will be strengthened with f_Y, the level of final output automation); (ii) the feedback loops induced through technological and economic channels as automation accelerates research (the next three terms, f_i scaled down by diminishing returns β_i, for each research sector).[2]

This condition by itself 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 & quantitative results. The explosion condition highlights the critical importance of measuring diminishing returns sector by sector. Bloom et al.

[1]For simplicity, the expression here assumes no "parallelization penalty"; (56) generalizes.

[2]In the case of general innovation, A, diminishing returns are modulated by the labor share, $1 - \alpha$.

(2020) estimate that in the economy as a whole, ideas become sharply harder to find, with $\beta_A = 3.1$. The hardware sector exhibits the same phenomenon – more and more researchers are required to maintain the pace of Moore's Law – but the magnitude is dramatically smaller: $\beta_H = 0.2$. Indeed, hardware shows the smallest degree of diminishing returns of any sector studied. Estimating diminishing returns in software has proven more challenging; we adopt a relatively conservative calibration of $\beta_S \approx 1$, drawing on Erdil, Besiroglu and Ho (2024)'s preferred estimate from chess engine data, while noting that other recent estimates point to lower values (Ho and Whitfill, 2025).

Applying this calibration to the explosion condition produces three analytic findings that only depend on the estimates of the diminishing returns to research and the economy-wide labor share. First, the threshold is reachable: 13% automation across all sectors is sufficient to push the economy into the explosive regime, and 17% suffices when only software and hardware research are automated. Second, hardware research is the dominant lever – because returns to research in hardware are roughly five times those in software and ten times those in aggregate TFP, automating one task in chip design moves the economy as much as five tasks in software or final-goods production. 20% automation of hardware alone is enough to cross the threshold. Third, software automation in isolation sits approximately at the knife-edge: under a fairly conservative calibration, fully automating software research without automating any other part of the economy just reaches the explosive growth threshold. A small push elsewhere is sufficient to tip the system.

We also simulate how quickly the singularity arrives once the threshold is crossed, an exercise which requires a calibration to current algorithmic and hardware progress (Epoch AI, 2026). In our baseline stylized simulation, an 'automation shock' involving full automation of software R&D and just 5% automation across the rest of the economy causes the singularity to arrive in roughly six years.

The role of bottlenecks. In our final section, we generalize our framework beyond the unitary elasticity of substitution case to allow for bottlenecks and show how *continually progressing* automation – dynamically evolving f terms – can overcome the constraining effects of such gross complementarity. We show that there is a race between the constraining effects of bottlenecks versus the amplifying effects of automation and analytically derive a sufficient condition on the speed of automation that restores the exact explosion dynamics from our baseline model.

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. The authors 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 and simulate their future impact.

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 degrees of diminishing returns, providing a unified condition for when growth becomes explosive. Second, we integrate technological heterogeneity with economic feedback loops, showing how output–technology complementarities amplify 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) model the growth implications of AI in a computational integrated assessment model. Our framework is simpler and sufficiently tractable to capture similar core dynamics analytically. This allows us to explicitly characterize feedback loops and their interactions and to derive an explicit explosion threshold.

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 assuming constant returns to scale on cross-sector innovation spillovers, ensuring balanced growth in an endogenous growth setting. By relaxing this assumption and introducing automation, we can examine how AI progress driven by software and hardware improvements may give rise to explosive growth.

Outline. Section 2 develops a set of simple models to illustrate the core economic forces at work in our setting, before sections 3 and 4 present the fully general network model and introduce automation. Section 5 presents our simple integrated AI economy, including the simulation in section 5.3; these can be read without the more general results. Section 6 shows how the results generalize to a world with bottlenecks, before section 7 concludes.

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 feedback loops (or amplifies existing spillovers).

2.1 Lesson one: Diminishing returns prevent growth explosions

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 to:

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

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:

$$Y_t = S_t L_t^{1-\alpha} \tag{3}$$

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.

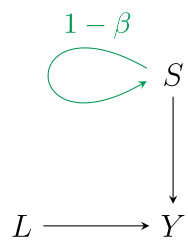

Figure 2: Recursive self-improvement explodes if and only if there are no diminishing returns: $\beta < 0$.

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; additionally, it is 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. Continuing to denote S as software productivity, denoting H as hardware quality, and dropping time subscripts to ease notation,[3]

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

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

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

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 $\phi_S = \phi_H = 0, \beta_H = 1$.

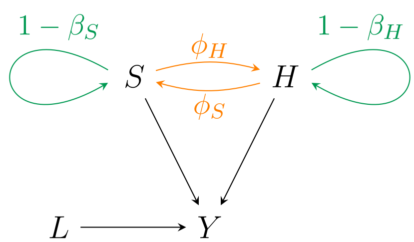

Figure 3: Innovation networks introduce spillovers, offsetting diminishing returns.

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.

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

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

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

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 $Y = K^\alpha$ 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 – $\alpha > 1$ – or in the language used here, if the economic 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:

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

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

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

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.

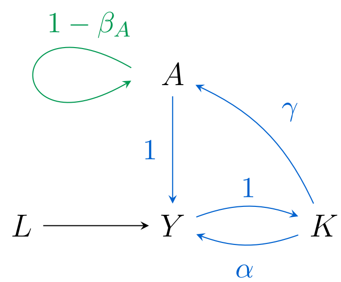

Figure 4: Economic feedback loops effectively offset diminishing returns.

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:

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

This condition is exactly analogous to the condition of lesson two, (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 ($\gamma > 0$), then explosive growth is more likely.

Notably, a standard calibration of (12) would imply a lack of explosive growth. Using $\beta_A = 3.1$ as the extent to which ideas are getting harder to find in the economy as a whole (Bloom et al., 2020), $1 - \alpha = 0.6$ as the labor share in production, and $\gamma = 0.1$ as the capital share in R&D for the economy as a whole (Besiroglu, Emery-Xu and Thompson, 2024), we find that the explosion condition is far from being met. However, as the next section shows, 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 in the simplest example tasks X_i are bundled into aggregate output Y via a Cobb-Douglas aggregate:[4]

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

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[5], 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)-(11), and consider changes in the

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

[5]Under an assumption that total capital K is sufficiently plentiful.

capital share in both production functions:

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

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

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

Here, $f_A \in [0,1]$ is the degree of automation in the ideas 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, makes clear how automation either strengthens existing feedback loops – or creates new ones which did not exist previously.

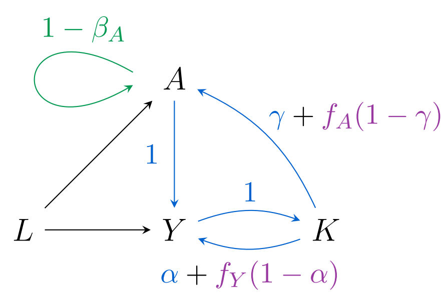

Figure 5: Automation introduces new feedback loops or strengthens existing ones.

For example, suppose initially capital was not used at all in production of ideas and no relevant tasks were automated: $\gamma = 0$ 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:

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

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

Using this last condition, we can now return to the boxed equation shown in the introduction, (1), to provide some intuition for its origin. Recall that condition:

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

Now consider the condition we derived in lesson four, (16). Our main application sets the initial capital share in research of zero, which would be equivalent to setting $\gamma = 0$.[6] Using this, the lesson-four condition (16) can easily be rewritten as:

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

Clearly, this condition matches the first and last terms of the boxed condition. 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 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 $\gamma > 0$ would introduce an additional economic feedback loop and thus 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 to 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:

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

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:

$$r_i = \frac{\lambda_i}{\beta_i} \tag{19}$$

$$s_{i,j} = \frac{\phi_{i,j}}{\beta_i} \qquad \text{for } i \neq j \tag{20}$$

The term r_i is a sector-specific measure of dynamic diminishing returns, capturing how technological progress responds to 'own-sector' research effort along the balanced growth path. 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:

$$\mathbf{F}^A_{i,j} = \begin{cases} s_{i,j}, & \text{if } i \neq j \\ 0, & \text{if } i = j \end{cases}. \tag{21}$$

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

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

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

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 paths*

$$\boxed{g_A^{BGP} = \mathbf{\Psi} \mathbf{r} g_E} \tag{22}$$

where $\mathbf{\Psi} = (\mathbf{I} - \mathbf{F})^{-1} \in \mathbb{R}^{N \times N}_{\geq 0}$ and $\mathbf{r} = diag(\ell_1/b_1, \ldots, \ell_N/b_N)$.

2. $\rho(\mathbf{F}) = 1$. *In the long run, the system 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.*

Proof. See Appendix A. □

This proposition says that key determinant of the long-run behavior of an innovation network like (17) is determined by the spectral radius (here, the largest eigenvalue) of the spillover matrix $\mathbf{F}^A$. In particular, the system exhibits explosive growth if and only if that spectral radius is greater than 1. This generalizes the single-sector explosion condition heuristically introduced in the introduction: $\frac{dy}{dt} = y^\beta$ for some $y \in \mathbb{R}_{++}$ explodes if and only if $\beta > 1$.

A complementary implication of proposition 1 is 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 $\mathbf{r}$ 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 supposing research inputs are proportional to population growth, we have the vector of balanced growth paths across technological sectors, assuming it exists:

$$g_A^{\text{BGP}} = \mathbf{rn} \ ,$$

where $\mathbf{n}$ is a column vector containing entries of the population growth rate, 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 diminishing 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 $s_{1,2}^{-1}$ and is $s_{1,2}$ for the $\hat{g}_{A_2}$ isoline.[8] Therefore, lines fail to cross whenever $s_{1,2}s_{2,1} > 1$. This is exactly the case where the largest eigenvalue of $\mathbf{F}^A$ exceeds one and there is no well-defined $\mathbf{\Psi}_A$.

Figure 6: Hyperbolic vs balanced growth in a two technology network

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

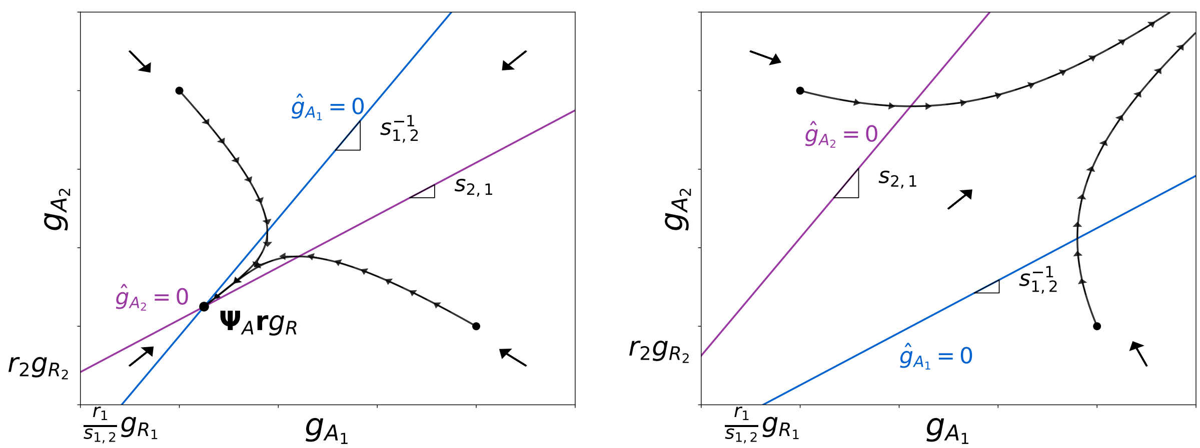

Note: Along colored lines the growth rate one of the technologies is constant (i.e., the growth-in-growth rate is zero). When these lines intersect the system exhibits stable balanced growth.

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 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*.

[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 20. That is, along the isolines we have

$$g_{A_1} = r_1 g_{R_1} + s_{1,2} g_{A_2} \qquad \text{and} \qquad g_{A_2} = r_2 g_{R_2} + s_{2,1} g_{A_1}$$

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 result in faster technological progress.

Specifically, consider the following system:

$$Y = \bar{A} K_Y^{\alpha} L_Y^{1-\alpha} \tag{23}$$

$$\dot{A}_i = v_i (K_i^{\gamma_i} L_i^{1-\gamma_i})^{\lambda_i} A_i^{1-\beta_i} \prod_{j \in I \setminus i} A_j^{\phi_{i,j}} \tag{24}$$

$$\bar{A} = \prod_{i \in I} A_i^{\tau_i} \quad \text{where} \quad \sum_{i \in I} \tau_i = 1 \tag{25}$$

$$\dot{K} = aY - \delta K \tag{26}$$

$$K = K_Y + \sum_{i \in I} K_i \tag{27}$$

$$L = L_Y + \sum_{i \in I} L_i \tag{28}$$

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 and population growth, n, is vectorized in $\mathbf{n}$.

Compared to the technology only system studies above, this system introduces a new differential equation for capital. This additional differential equation requires some special treatment given the linear drag imposed by depreciation, which is clearly not a feature of the technology differential equations. Fortunately, Jones (2026) demonstrates in a simpler technology-capital system the level of depreciation will not impact whether this system exhibits explosive growth. Ultimately we demonstrate that this observation holds in this growth system, and as a result we can partition the growth environment almost identically to in Proposition 1.

Proposition 1.* (Economic and Technological Feedback). *Take the system of equations described in* (23)-(28)*. Further, define the economic feedback matrix* $\mathbf{F}^Y := \frac{1}{1-\alpha}[\mathbf{r}\gamma]\tau'$*, where* τ *is the vector of exponents from equation* (25)*, and* $\mathbf{r} = diag(r_1, \ldots r_N)$*.*

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

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

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* $N \times N$ *technological and economic feedback matrices, with* $\mathbf{F}^A$ *and* $\mathbf{r}$ *defined above (equations (21) and (19)).*

2. $\rho(\mathbf{F}^A + \mathbf{F}^Y) = 1$. *In the long run, the system is bounded by double-exponential growth, and exponential in the purely endogenous growth case* ($n = 0$).

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

Proof. See Appendix A. □

The key difference between this result and Proposition 1 is that we can just replace the instances of feedback matrix $\mathbf{F}$ with two separate feedback matrices $\mathbf{F}^A$ and $\mathbf{F}^Y$. Here the technology feedback matrix, $\mathbf{F}^A$, is just that defined above. The economic feedback matrix, $\mathbf{F}^Y$, captures the degree to which specific technologies increase output and that output subsequently increases technological progress along the balanced growth path. The vector of terms, $\mathbf{r}\gamma$, 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 (24) we can arrive at

$$g_{A_i}^{\text{BGP}} = r_i\gamma_i g_Y^{\text{BGP}} + \sum_{j\in I\setminus i} s_{i,j} g_{A_j}^{\text{BGP}} + r_i(1-\gamma_i)n. \tag{30}$$

Further, the vector of terms $\frac{1}{1-\alpha}\tau$ mediates the effect of technological progress on production; on the balanced growth path, from equation (23) we can arrive at

$$g_Y^{\text{BGP}} = \frac{1}{1-\alpha}\tau \cdot g_A^{\text{BGP}} + n. \tag{31}$$

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

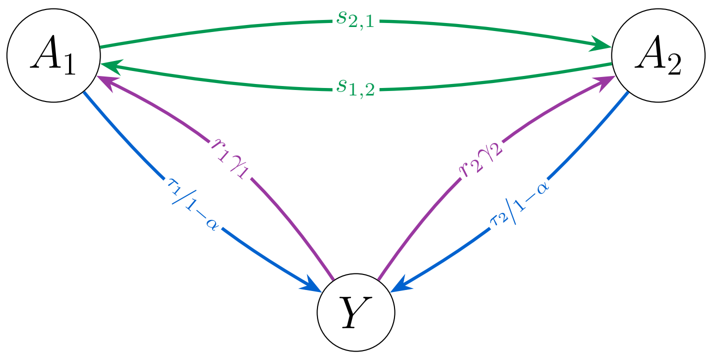

Note: Summarizing entries in $\mathbf{\Psi}_{A,Y}$, green lines represent the technology spillovers, $s_{i,j}$, that make up entries of $\mathbf{F}^A$; purple lines represent output-technology feedback mediated by the vector $\mathbf{r}\gamma$ from $\mathbf{F}^Y$; blue lines represent the technology-output feedback mediated by the vector $\tau(1-\alpha)^{-1}$ from $\mathbf{F}^Y$.

The definition of the Leontief inverse, $\mathbf{\Psi}_{A,Y}$, provides a clear decomposition of the effects of technological and economic feedback loops to inform the balanced growth path of this economy. Just as with the balanced growth path under technological feedback loops alone, we can recover the standard balanced growth path from Jones (1995) if $\mathbf{\Psi}_{A,Y} = \mathbf{I}$. 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 and to total factor productivity, so $\mathbf{F}^Y_{i,j} = 0$ for all i and j (that is, if $\gamma_i > 0$ then require $\tau_i = 0$ and if $\tau_i > 0$ require $\gamma_i = 0$).

In this model we assume that savings rates are constant (but may be arbitrary levels). Of course, optimal savings paths will respond endogenously to changes 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, a decline in savings 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

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 p_i is a productivity constant and

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

where task q is performed either with human labor $L_{i,q}$ or AI labor $C_{i,q}$. Let f_i denote the fraction of tasks in sector i which are automatable by AI; assuming that compute capital is sufficiently abundant, all tasks which are automatable are automated. Then, it is optimal to spread AI labor equally across automated tasks and to spread human labor equally across non-automated tasks. Setting aside integer constraints, effective labor is thus given by:

$$\hat{L}_i(C_i, L_i) = C_i^{f_i} L_i^{1-f_i} \tilde{p}_i(f_i) \tag{32}$$

where $\tilde{p}_i(f_i) \equiv \frac{p_i}{f_i^{f_i}(1-f_i)^{1-f_i}}$, which is an unimportant constant for a given level of automation and does not affect explosive dynamics. Importantly, we here treat the level of automation f_i as an exogenous constant to evaluate the growth implications of a given level of task automation.

4.2 AI-driven automation

Sector-i AI labor, C_i, is determined by both hardware used in that sector and the level of software capabilities in the economy. We present the notation and then explain.

AI labor in sector i is:

$$C_i = \underbrace{K_{H \to i}}_{\text{rival}} \times \underbrace{\prod_{j \in \bar{I}} A_j^{\sigma_j}}_{\text{non-rival}} \tag{33}$$

Hardware is a particular kind of capital – inference compute – required to perform AI labor. We denote the quantity used in sector i as $K_{H \to i}$. Note this is a rivalrous good: use in one sector precludes use in another sector. Total inference compute across the economy is denoted K_H.

The level of software capabilities in the economy – also frequently termed "algorithmic progress" – is nonrivalrous and determined by a composite of ideas from different research sectors, $\prod_{j \in \bar{I}} A_j^{\sigma_j}$. We allow $\sigma_j = 0$ if a sector does not contribute to software progress, and we do not necessarily assume that $\sum_{i \in \bar{I}} \sigma_j = 1$ since standard replication arguments for constant returns to scale do not apply to non-rivalrous technologies. Examples of algorithmic progress include the original idea of a neural network; the modern deep learning paradigm; or the transformer.

We take software capabilities to raise the productivity of a given quantity of compute. The choice to model AI progress as arising from both inference compute scaling as well as software progress is inspired by the crucial empirical fact that AI progress has been driven by both algorithmic advances and also from increasing compute inputs (Ho et al., 2024).

4.3 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 (32). Therefore the technology laws of motion become

$$\begin{aligned}\dot{A}_i &= \nu_i \hat{L}_i^{\lambda_i} A_i^{1-\beta_i} \prod_{j \in I \setminus i} A_j^{\phi_{i,j}} \\ &= \nu_i \tilde{p}_i(f_i) C_i^{f_i \lambda_i} L_i^{(1-f_i)\lambda_i} A_i^{1-\beta_i} \prod_{j \in I \setminus i} A_j^{\phi_{i,j}} \\ &\propto \underbrace{A_i^{1-\beta_i}}_{\text{diminishing returns}} \times \underbrace{[K_{H \to i}^{f_i} L_i^{(1-f_i)}]^{\lambda_i}}_{\text{rival inputs}} \times \underbrace{\prod_{j \in I} A_j^{f_i \lambda_i \sigma_j}}_{\text{AI feedback}} \times \underbrace{\prod_{j \in I \setminus i} A_j^{\phi_{i,j}}}_{\text{Direct spillovers}} \end{aligned} \tag{34}$$

As a baseline, we assume that direct technological spillovers are non-existent ($\phi = 0$). In this case, we recover sectoral spillovers through research 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 model here to exactly match the earlier model, as summarized 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	$1 - f_i$	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 (by allowing $\phi_{i,j} > 0$), 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 (22) and denoting automation-adjusted parameters with hats, we have

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

where

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

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. And as above, $\hat{\mathbf{r}} = diag(\hat{r}_1, \ldots, \hat{r}_N)$.

We can apply Proposition 1 directly to this system. This yields:[9]

Corollary 1. *The automation-calibrated technology system (equation (34)) explodes in finite time iff*

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

Proof. See Appendix A. □

4.4 Hyperbolic growth: Automated production and research

Now we extend the AI-induced feedback loops captured in equation (34) 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.

$$Y = \bar{A} {K_Y}^{\alpha} \hat{L}_Y^{1-\alpha} \tag{35}$$

$$\bar{A} = \prod_{i \in I} A_i^{\tau_i} \quad \text{where} \quad \sum_{i \in I} \tau_i = 1 \tag{36}$$

$$\dot{A}_i \propto \hat{L}_i^{\lambda_i} A_i^{1-\beta_i} \tag{37}$$

$$\hat{L}_i \propto L_i^{1-f_i} C_i^{f_i} \quad \text{where} \quad C_i = K_{H \to i} \times \prod_{j \in I} A_j^{\sigma_j} \tag{38}$$

$$\dot{K}_Y = a_Y Y - \delta_Y K_Y \tag{39}$$

$$\dot{K}_H = a_H Y - \delta_H K_H \tag{40}$$

$$K_H = K_{H \to Y} + \sum_{i \in I} K_{H \to i} \tag{41}$$

$$L = L_Y + \sum_{i \in I} L_i \tag{42}$$

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

$$g_A^{\text{BGP}} = \hat{\mathbf{\Psi}}_{A,Y} \hat{\mathbf{r}} \mathbf{n} \tag{43}$$

[9] We note that this result takes terms $\hat{r}$ as an input; however, the final form is written in terms of the unadjusted returns to research, $r_i = \lambda_i / \beta_i$.

where

$$\hat{\boldsymbol{\Psi}}_{A,Y} = (\mathbf{I} - [\hat{\mathbf{F}}^A + \hat{\mathbf{F}}^Y])^{-1} \tag{44}$$

$$\hat{\mathbf{F}}^Y := \underbrace{[\hat{\mathbf{r}} f]}_{dg_A^{\text{BGP}}/dg_Y^{\text{BGP}} \text{ from } 37} \underbrace{\left[\frac{1}{1-f_Y}[\frac{1}{1-\alpha}\tau + f_Y \sigma]\right]'}_{dg_Y^{\text{BGP}}/dg_A^{\text{BGP}} \text{ from } 35} \tag{45}$$

Here, $\hat{\mathbf{r}}$ and $\hat{\mathbf{F}}^A$ are defined as above, and f is the vector of automation shares in each technology sector.

In table 2 we present how the core dynamics of the general R&D and production network model described by equations (23)-(28) can be calibrated to the automated economy model in equations (35)-(42) 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: increasing the capital share of production by $f_Y(1-\alpha)$; and increasing the contributions of each technology to the production of final goods by $f_Y \sigma$.

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

General Model Input	**AI-Parameterization**	**Description**
α	$\alpha + f_Y(1-\alpha)$	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	$1 - f_i$	Non-automated research task share
λ_i	λ_i	Parallelization penalty

Note: This table summarizes how one can re-parameterize the general model (equations (23)-(28)) 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.

Applying Proposition 1 to the balanced growth condition from equation (43), we have the following result:

Corollary 2. *With economic feedback loops, the automation-calibrated growth model (described in equations (35) - (42)) explodes in finite time iff*

$$f_Y + \sum_{i \in I} r_i f_i \left(\frac{\tau_i}{1-\alpha} + \sigma_i\right) > 1 \ . \tag{46}$$

Proof. See Appendix A. □

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

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

Above we have introduced a general model of networked growth, where we 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:

$$Y = A\hat{L}_Y^{1-\alpha} {K_Y}^{\alpha} \tag{47}$$

$$\hat{L}_i \propto L_i^{1-f_i} C_i^{f_i} \quad \text{where} \quad C_i = K_{H\to i} S \tag{48}$$

$$\dot{S} \propto \hat{L}_{\dot{S}}^{\lambda_S} S^{1-\beta_S} \tag{49}$$

$$\dot{H} \propto \hat{L}_{\dot{H}}^{\lambda_H} H^{1-\beta_H} \tag{50}$$

$$\dot{A} \propto \hat{L}_{\dot{A}}^{\lambda_A} A^{1-\beta_A} \tag{51}$$

$$\dot{K}_H = a_H H Y - \delta_H K_H \tag{52}$$

$$\dot{K}_Y = a_Y Y - \delta_Y K_Y \tag{53}$$

$$K_H = K_{H\to Y} + K_{H\to \dot{H}} + K_{H\to \dot{S}} + K_{H\to \dot{A}} \tag{54}$$

$$L = L_Y + L_{\dot{H}} + L_{\dot{S}} + L_{\dot{A}} \tag{55}$$

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 inference of AI models. This channel from output to compute is amplified through hardware quality progress in (52), which is a form of investment-specific technical change (more inference compute can be created from a given amount of raw investment over time). 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 4.4, we can find the conditions for balanced and explosive growth by recognizing from equations (52) and (53) that a balanced growth path requires $g_{K_Y} = g_{K_H} - g_H = g_Y$. 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 4.4. 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 (49)-(55) is a three technology version of the general model from equations (35)-(42), where: technological contributions to output are dictated by $\tau_S = \tau_H = 0$ and $\tau_A = 1$; and technological contributions to AI progress are dictated by $\sigma_S = \sigma_H = 1$ and $\sigma_A = 0$.[10]

Given that the AI-economic model developed here is a specific case of the general model, we can simply calibrate Corollary 2 to derive the hyperbolic growth condition that was first presented in equation (1) in the introduction:

Corollary 3. *The AI-economic model (described in equations (47)-(55)) explodes in finite time iff*

$$f_Y + f_S r_S + f_H r_H + f_A \frac{r_A}{1-\alpha} > 1 \tag{56}$$

5.2 Calibration: Explosive vs balanced growth

We now turn to calibrating the relevant terms from equation (56). In table 3 we report estimates of diminishing returns in software, hardware and aggregate TFP. Both software and (in particular) hardware research face significantly less diminishing returns than aggregate TFP. That is, ideas are getting harder to find in all sectors, but less so in software and particularly hardware. We discuss limitations of this calibration in section 5.4.

Table 3: Parameter estimates

Term	Parameter	Estimate	Source
Labor share	$1-\alpha$	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)

[10] The only difference between this calibrated model and the general model from section 4.4 is that one of the technologies – hardware quality – scales compute accumulation rather than AI capabilities directly. This does not affect the balanced growth path calculation.

Table 4 presents the calibrated explosion conditions, i.e. from calibrating (56) using table 3.

Table 4: Applying historical estimates of diminishing returns to singularity conditions

Returns to Capital + Labor	Calibrated Explosion Condition
Constant (no fixed factor)	$\sim 1f_S + 5f_H + 0.53f_A + f_Y > 1$
10% fixed factor	$\sim 1f_S + 5f_H + 0.5f_A + 0.84f_Y > 1$

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.[11] Appendix B provides the analytical balanced and hyperbolic growth conditions with a fixed factor.

Turning to table 5, each row considers automating a different set of sectors. Assuming an equal level of automation in each sector under consideration, and zero automation in other sectors, column two solves for how much automation is required to achieve a doubling of output growth on the balanced growth path. Column three, 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 most of the automation necessary to achieve fully explosive growth. That is, the balanced growth path is initially very slow to respond to changes in automation, and then changes 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

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

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.

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 viewpoint of the entire human history the current 2% annual growth is part of a long-run superexponential trend (Hanson, 2000; Roodman, 2020).

That said, empirically the recent growth rates of productivity in software and hardware have been so extraordinarily fast, and so it is also plausible that the transition to a new balanced growth path or hyperbolic acceleration happens extremely quickly. We turn to quantifying these considerations in the next section.

5.3 Simulation: The effect of an automation shock over time

We now simulate an entire timepath for technology and output following a one-time increase in the share of automated tasks, and (when the explosion condition is satisfied) calculate the singularity date implied by the model. Specifically, starting from zero automation, we simulate the model after a jump in sectoral automation rates: for $i \in \{S, H, A, Y\}$,

$$f_{i,t} = \begin{cases} 0 & \text{if} \quad t < 0 \\ f_i & \text{if} \quad t \geq 0 \end{cases} \tag{57}$$

Calibration. The simulation require introducing additional quantitative assumptions beyond those in the previous subsections.

We assume the economy starts on a balanced growth path. To calibrate the research sectors, we use technological growth rates by sector, (43). From this equation, we can back out researcher growth in each sector (n_H and n_S) by combining the ideas-getting-harder-to-find parameters in table 3 together with measured, current technological growth rates in software and hardware reported by Epoch AI (2026):

$$\boxed{g_{S,2026} = 110\% \quad \text{and} \quad g_{H,2026} = 37\%} \tag{58}$$

Together these implies that the number of researchers in the hardware sector has been growing at $n_H = 7.4\%$ and in the software sector at $n_S = 110\%$ per year.[12] Additionally, Bloom et al. (2020) report that quality-adjusted workers in aggregate R&D have been increasing at a rate of 4% annually.

In the output sector, we assume that the growth rate of workers is 1%.[13]

In the task model, an automation shock produces two growth effects: reallocating human workers into a narrower range of tasks, and increasing the elasticity of effective labor toward an endogenously accumulating factor. The first causes an immediate jump

[12] Our calibration of hardware researcher growth, $n_H = 7.4\%$, happens to be identical to the estimate of Bloom et al. (2020) of growth in research inputs in the hardware sector. Our calibration of software researcher growth, $n_S = 110\%$, is lower than, but of the same order of magnitude as, the estimate by Epoch AI (2026) of growth in AI lab headcount of roughly 140% per year since ChatGPT.

[13] In Appendix C we also consider a version of the exercise where the growth rates of research inputs are just chosen to match an aggregate population growth rate of 1%. Since the zero-automation balanced growth path in each technology sector is $r_i \times n_i$, this just means the technology balanced growth paths are 0.3%, 1%, and 5% for TFP, software and hardware – clearly very different from the rates from (58). This difference is consequential for singularity timing when the explosion condition is satisfied.

in growth rates; the second shapes the long-run path of variables. We focus on the second and therefore assume no labor reallocation after the shock. Mathematically, this amounts to setting $\tilde{p}(f_i) = 1$ for all i in the effective labor function $\hat{L}_i = \tilde{p}(f_i) L_i^{1-f_i} C_i^{f_i}$.

In addition to these key calibrations, we assume:

- Depreciation is 7% for conventional capital and 20% for computing hardware. Along with the balanced growth path assumption, this pins down investments rates $(a_Y, a_H) = (0.41, 0.25)$. While these rates are unrealistic, this is an implication of the assumption that we are currently on a balanced growth path. The lower balanced growth path considered in Appendix C allows for more realistic investment rates (and correspondingly slower growth accelerations).

- Along the no-automation balanced growth path, the conventional capital to output ratio is 4. This implies that on the balanced growth path the compute to output ratio is 0.41.[14]

- Compute stocks at $t = 0$ are large enough such that the effective labor is unchanged by the shock. Therefore,

$$L_{i,0}^{1-f_i} C_{i,0}^{f_i} = L_{i,0} \implies C_{i,0} = L_{i,0} \quad \text{for any } f_i$$

  implying that technology growth rates do not jump at $t = 0$, and all acceleration comes from accumulation and technological progress.

- The share of compute allocated to each sector is equal to the share of human labor allocated to each sector ($\ell_i = \kappa_i$). This implies the allocation of effective labor across tasks is unchanged after the automation shock.

Note, under these assumptions (in particular, $\ell_i = \kappa_i$) we do not need to assume the levels of shares (ℓ_i, κ_i). This is because the labor shares are jointly determined with the coefficient that pins down growth levels.[15]

[14] Even though compute is not economically useful in this model before $t = 0$ we assume it is accumulated regardless so that effective labor is unchanged after the automation shock.

[15] To illustrate, suppose

$$g_{S,t} = v_S \hat{L}_{S,t}^{\lambda_S} S_t^{-\beta_S}$$

then for $t < 0$ (along the no-automation balanced growth path)

$$g_S^{BGP} = v_S (\ell_S L_t)^{\lambda_S} S_t^{-\beta_S} \implies v_S (\ell_S)^{\lambda_S} = g_S^{BGP} L_t^{-\lambda_S} S_t^{\beta_S}$$

Figure 8: Automation shock experiment: Uniform automation across sectors

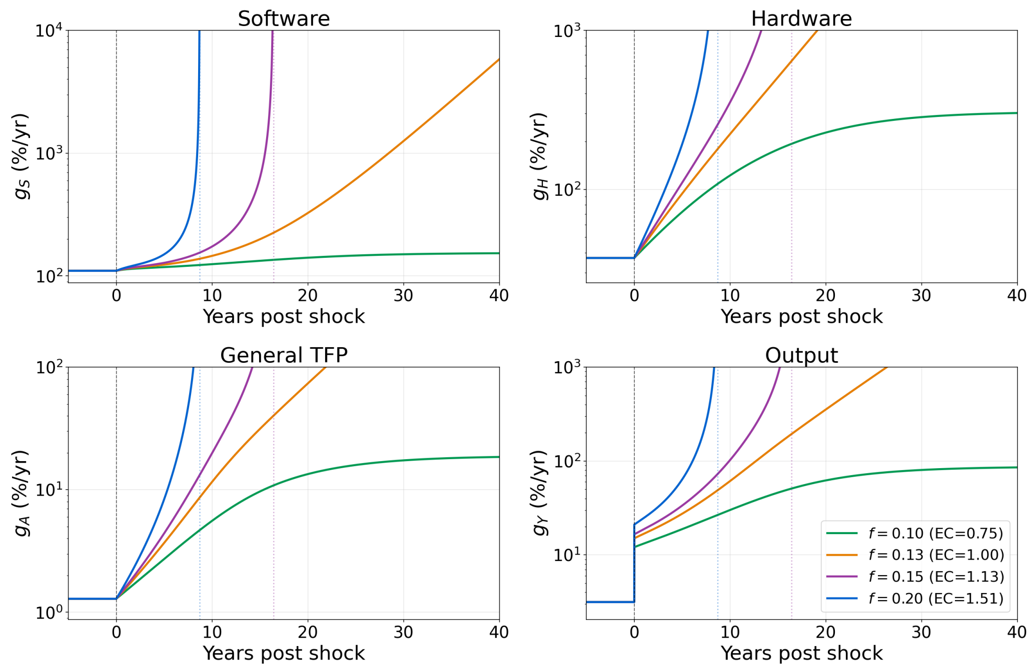

Note: Response of growth rates in software (g_S), hardware (g_H), general TFP (g_A), and output (g_Y) to a one-time, permanent increase in automation at $t = 0$, starting from a balanced growth path. A uniform automation shock is applied across all sectors ($f_S = f_H = f_A = f_Y = f$). The legend reports the "explosion condition" value $EC \equiv f_Y + f_S r_S + f_H r_H + f_A \frac{r_A}{1-\alpha}$, with r and α values from Table 3. Per Proposition 1, $EC > 1$ yields hyperbolic growth (finite-time singularity), $EC = 1$ yields double-exponential growth, and $EC < 1$ yields balanced exponential growth. Here, 13% uniform automation is the knife-edge case ($EC = 1$); 15% and 20% cross into the explosive regime, producing singularities in roughly 16 and 9 years respectively.

Figure 9: Automation shock experiment: Software-focused automation

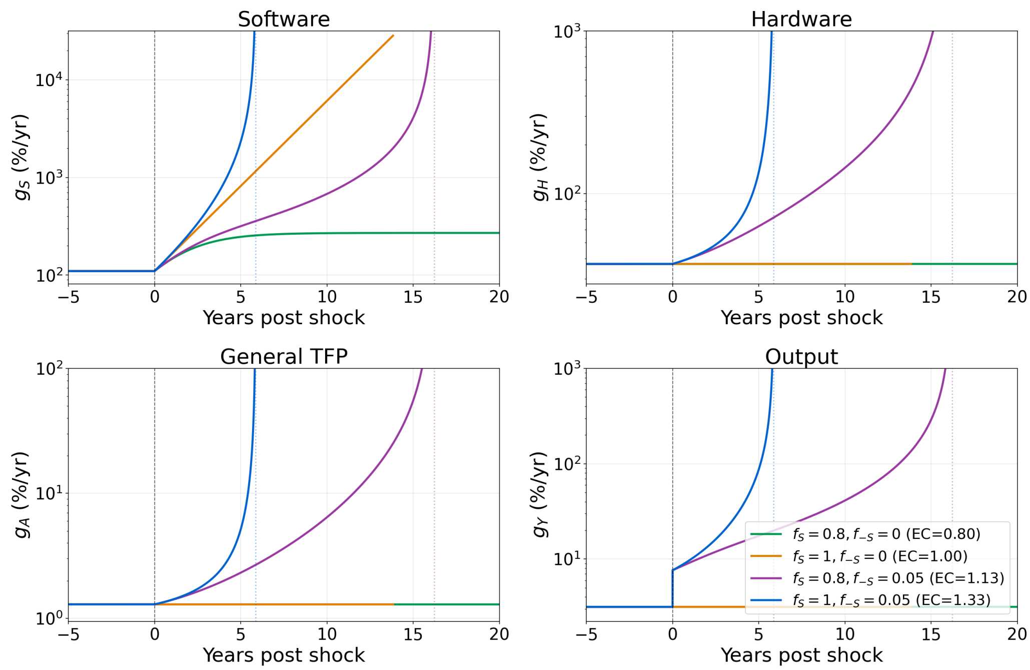

Note: Response of growth rates in software (g_S), hardware (g_H), general TFP (g_A), and output (g_Y) to a one-time, permanent increase in automation at $t = 0$, starting from a balanced growth path. A large shock is applied to software research automation ($f_S \in \{0.8, 1\}$) while automation in all other sectors is kept low ($f_{-S} \in \{0, 0.05\}$). The legend reports the "explosion condition" value $EC \equiv f_Y + f_S r_S + f_H r_H + f_A \frac{r_A}{1-\alpha}$, with r and α values from table 3. Per Proposition 1, $EC > 1$ yields hyperbolic growth (finite-time singularity), $EC = 1$ yields double-exponential growth, and $EC < 1$ yields balanced exponential growth. Fully automating software alone ($f_S = 1$, $f_{-S} = 0$) is exactly the knife-edge case, illustrating that software automation in isolation is insufficient for hyperbolic growth without at least some automation elsewhere.

Simulation results. In figures 8 and 9, we illustrate the growth rate of our four main variables (software, hardware, general TFP, and output) over time following the jump in automation from $f = 0\%$ to a higher level at time zero. Figure 8 considers a uniform automation level ($f_i = f$ for all i), where f ranges from 10% to 20%. Figure 9 considers the case of high software automation ($f_S \in \{0.8, 1\}$), and lower automation in the rest of the economy ($f_{-S} \in \{0, 0.05\}$). In the legend of these figures we report the level of the left hand side of the explosion condition ('EC') in equation (56) that determines whether there is hyperbolic (EC > 1), double exponential (EC $= 1$), or balanced growth (EC < 1) as stated in Proposition 1.

We can see that in both the uniform automation and the high software automation cases we illustrate a double exponential growth path example: $f = 0.13$; and $f_S = 1$ and $f_{-S} = 0$ respectively. In the uniform automation case, higher automation than 13% implies hyperbolic growth, with singularity dates of nine and 16 years for $f = 0.15$ and $f = 0.2$; while less than 13% automation results in the system converging back to a balanced growth path, for $f = 0.1$ this growth path is still very large, with output doubling yearly. In the uniform automation case there is an initial jump in the growth rate of output, since the effective labor remains unchanged, but the capital share is increased.[16] In the high software automation case, we only see changes in the growth of output, hardware quality and general TFP when these sectors can be automated to some degree; in our case, $f_{-S} = 0.05$. In those cases there is sufficient automation across the economy for hyperbolic growth, with singularity dates 6 and 16 years after the automation shock depending on whether there is full or partial automation ($f_S = 1$ or $f = 0.8$).

In figure 10, we plot the singularity date from the above calibration while varying the level of automation and the sectors in which automation occurs.[17] With full au-

where the RHS is constant along the balanced growth path. Then, for $t \geq 0$ we have

$$\begin{aligned} g_{S,t} &= v_S(\ell_S L_t)^{\lambda_S(1-f)}\left(\kappa_S K_{H,t} S_t\right)^{\lambda_S f} S_t^{-\beta_S} \\ &= v_S \ell_S^{\lambda_S} L_t^{\lambda_S(1-f)} (K_{H,t} S_t)^{\lambda_S f} S_t^{-\beta_S} \qquad (59) \\ &= g_S^{BGP} (K_{H,t} S_t / L_t)^{\lambda_S f} \qquad (60) \end{aligned}$$

where (59) comes from $\ell_i = \kappa_i$ and (60) comes from the solution for $v_S \ell_S^{\lambda_S}$. Therefore, we can see that the growth path is pinned down entirely by the initial balanced growth path ($r_S n$) and the paths of aggregates, rather than allocations. This can be repeated for all technologies.

[16] Using the balanced growth paths of g_{K_H}, g_S, g_H and g_A and the equation for g_Y, one can show that the size of the jump in output is equal to $(1-\alpha) f (r_S n_S + r_H n_H + \frac{r_A n_A}{1-\alpha})$.

[17] Computationally, we approximate the singularity by testing when growth in the system exceeds 200% per year.

Figure 10: Automation shock experiment: Singularity dates over shock size

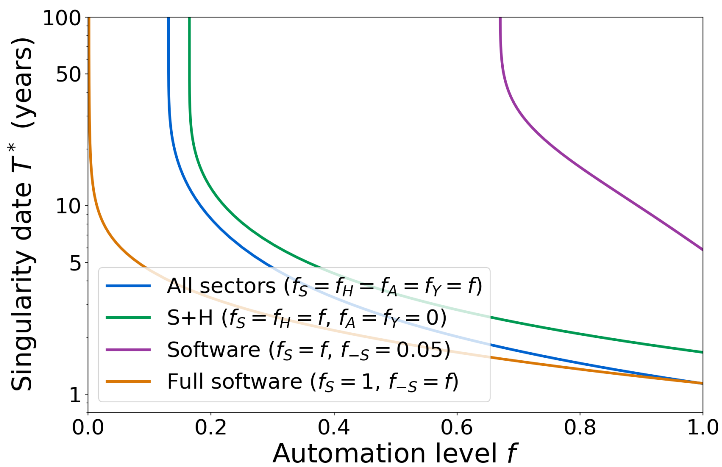

tomation across all sectors, the singularity arrives just one year after the shock, and later as fewer sectors are affected. As shown in table 5, the singularity date asymptotes at the automation thresholds required for explosive growth: 13% when all sectors are automated, 17% when only software and hardware sectors are automated, 66% when a software automation shock occurs while other sectors remain at 5% automation, and 0% when software is fully automated and the shock size varies across other sectors. This figure also illustrates that under a large automation shock, the growth model predicts a singularity very quickly. It is in these very high-growth scenarios that bottlenecks—ignored in this section—would become a particularly significant drag on growth. We turn to analyzing the role of bottlenecks in the following section.

This automation shock experiment is illustrative not only because we assume no bottlenecks, which are discussed further in the next section, but also because we hold the share of resources devoted to hardware and software research fixed at recent levels. Other work has endogenized these shares (Erdil et al., 2025); but here we maintain the simpler calibration to isolate the feedback effect of setting $f > 0$.

5.4 Limitations of calibration and simulation

There are a number of limitations to the estimates used to calibrate our model.

First, is that diminishing returns 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, in part because of challenges associated with measuring research inputs. Ho and Whitfill (2025) attempt to make this calculation directly, finding r_S 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 or 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 effective strengthening of the ideas-getting-harder-to-find effect as the average quality of researchers 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 using the aggregate economy estimate.

Fourth, our baseline environment (sections 5.2 and 5.3) assumes conventional capital is not a research input, so automation creates *new* feedback loops between output and technology. An alternative specification redefines effective R&D labor as

$$\hat{L}_i \propto K_i^{\gamma_i} (L_i^{1-f_i} C_i^{f_i})^{1-\gamma_i} \tag{61}$$

where γ_i is the elasticity of sector i research input with respect to conventional capital. Here, economic feedback loops exist even without automation ($f_i = 0$), and the techniques of section 3.2 can derive explosive and balanced growth conditions as a function of the research capital shares γ_i. In this case, adding automation ($f_i > 0$) on top of conventional capital in research would require less automation than reported in table 5 to trigger explosive growth, assuming the economy begins on a balanced growth path.

6 The role of bottlenecks

This section discusses (i) how to think about bottlenecks conceptually, (ii) the key economic object which determines whether a growth explosion is possible, and (iii) a condition under which the above explosion results continue to hold in the presence of bottlenecks.

The model in section 5 studies an economy without any bottlenecks. Formally, this is evident in (48), where effective labor is a Cobb-Douglas aggregate of human labor and AI labor, with $\hat{L} \propto L^{1-f}C^{f}$. That is, the elasticity of substitution between human labor and AI labor is one, so that an infinite amount of either input is sufficient for infinite output. An elasticity less than one would imply that output is bottlenecked by the slower-growing input – in particular, no amount of AI labor would compensate for a shortage of human labor.

Thus, as is well-known, bottlenecked production technology will never allow for an economic singularity except under full automation. In particular, there is no level for the AI-related diminishing return parameters β_S and β_H in the idea production functions (49)-(50) which would allow for a singularity. Even increasing returns, $\beta_i < 0$, would be insufficient. That's because, as the level of AI labor C skyrockets to infinity, the level of effective labor $\hat{L}$ is still bottlenecked by the amount of human labor. In the limit, $\hat{L} \propto L$.[18]

Of course, even if the elasticity of substitution were unity "in reality", so that singularities were possible as in the model, then reality could not feature a *true* singularity: to our best understanding, we exist in a finite bounded lightcone (Sandberg and Manheim, 2021). Thus, in *any* case, the model here should be understood like any other economic model: an approximation to reality, with a bounded domain of applicability, aiming to serve as a powerful intuition pump (Krugman, 1998). In *any* case, the conditions derived here could only possibly delineate a temporary period of 'explosive' growth – even if the 'temporary' period were a very long one.

[18]Likewise, no amount of spillovers in an innovation network would be enough to allow for an economic explosion, in the presence of these bottlenecks to automation.

Introducing bottlenecks. It is nonetheless interesting to think about how the presence of explicit bottlenecks affects our application in section 5. Allowing for bottlenecks would amount to replacing the aggregator for effective labor (48) with the following CES aggregator:

$$\hat{L}_i \propto \left((1-f_i)^{\frac{1}{\theta}} L_i^{\frac{\theta-1}{\theta}} + f_i^{\frac{1}{\theta}} C_i^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}} \tag{62}$$

with $\theta < 1$ allowing for bottlenecks. Meanwhile, $\theta \to 1$ recovers our baseline specification.

The key insight is that the important object is the *elasticity of effective labor to AI*:

$$\eta_i \equiv \frac{\partial \log \hat{L}_i}{\partial \log C_i} \tag{63}$$

In the Cobb-Douglas case, this elasticity is simply $\eta_i = f_i$.

In the CES case this elasticity becomes slightly more complicated. Denote the ratio of non-automated to automated tasks as:

$$x_i \equiv \frac{1-f_i}{f_i}.$$

Then the elasticity of effective labor to AI in the presence of bottlenecks is:

$$\eta_i = \frac{1}{1 + x_i^{\frac{1}{\theta}} (C_i/L_i)^{\frac{1-\theta}{\theta}}}. \tag{64}$$

This elasticity is now sensitive to not only the share of automated tasks f_i but also the AI-human labor ratio, C_i/L_i. As a result, as AI labor grows faster than does human labor, this elasticity declines – and so output becomes less and less responsive to increases in the AI labor supply. In the limit as $C_i/L_i \to \infty$, for $\theta < 1$, $\eta_i \to 0$.

In the bottlenecked economy, the relevant explosion condition is obtained by replacing the benchmark automation terms with the corresponding CES elasticities of (64). With the declining elasticity as C_i/L_i rises, growth in effective labor tapers off, thus preventing a growth explosion.

A sufficient condition for preserving the explosion threshold. We now introduce a sufficient condition under which our baseline results go through even under $\theta < 1$. The key change is that, rather than treating the level of automation f as a fixed constant, we suppose it evolves dynamically. Continual advancement in f then prevents the decline in the key elasticity η.

A condition on *the speed of automation* that is sufficient to prevent the elasticity from collapsing is the following:

$$g_{x,i} \leq -(1-\theta) g_{C_i/L_i} \tag{65}$$

where $g_{x,i} \equiv \frac{d \log x_i}{dt} = -\frac{g_{f,i}}{1-f_i}$ is the rate at which the share of tasks that are non-automated declines in sector i. As long as automation progresses at a speed no less than this threshold, the elasticity (64) will not fall.

Thus, if the relevant CES elasticities satisfy the critical threshold at some date, and if condition (65) holds thereafter, then the bottlenecked economy remains on the explosive side of the benchmark threshold.

We state these arguments formally in corollary 4.

Corollary 4. *Consider the bottlenecked AI-economic model described by equations (47), (62), and (49)-(55). Suppose the sectoral labor elasticities, $\eta_Y, \eta_S, \eta_H, \eta_A$, at some time T satisfy the non-bottlenecked threshold condition*

$$\eta_Y(T) + \eta_S(T) r_S + \eta_H(T) r_H + \eta_A(T) \frac{r_A}{1-\alpha} > 1 \tag{66}$$

If the speed of automation in each sector satisfies (65) for all $t \geq T$, then economic output explodes in finite time.

This result is simply our prior threshold condition (56), but replacing the f_i automation terms with the more general $\eta_{i,t}$ elasticities, together with a requirement that guarantees the elasticity (64) is non-decreasing over time. The CES labor aggregator (62) then admits a uniform lower bound of Cobb-Douglas form with fixed exponents $\eta_i(T)$. The bottlenecked economy therefore dominates a Cobb-Douglas comparison system with automation shares f_i set equal to $\eta_i(T)$; since that comparison system satisfies the threshold in Corollary 3, it explodes in finite time, and so does the CES economy.

Further characterization of the sufficient condition. In many models with an endogenous automation frontier, f (and hence x) is explicitly a function of the capital-to-labor ratio.[19] Suppose x is a function of C/L with an elasticity denoted $\varepsilon \equiv \frac{\partial \log x(C/L)}{\partial \log C/L}$.

[19]Where, for ease of notation, we now drop sectoral i subscripts.

Then $g_x = \varepsilon g_{C/L}$, so condition (65) can be rewritten as:

$$-\varepsilon \geq 1 - \theta \tag{67}$$

That is, consider a 1% increase in the AI-human labor ratio, C/L. This causes a $-\varepsilon\%$ fall in x, and approximately the same increase in f for $f \approx 1$. Equation (67) states that in response to this 1% increase we require that $(1-\theta)\%$ of the remaining tasks are automated. Note that the more complementary is the aggregation ($\theta \to 0$), the faster that tasks must be automated to sustain the required output elasticity for maintaining explosive growth.

Two remarks are worth noting. First, if $C/L \to \infty$, then under CES aggregation $\eta \to 0$ unless $f \to 1$, so eventual full automation is required to keep AI's output elasticity bounded away from zero. Second, if some tasks are never automatable, then the elasticity condition will eventually fail, even if the CES economy initially shadows the explosive benchmark.

Summary. Together, these points indicate that the possibility of task complementarities does not invalidate the conditions for explosive growth derived above. Our main results consider the Cobb-Douglas case because it has the straightforward property that the elasticity of the aggregator of AI-equivalent labor and human labor is just the share of tasks that can be automated. Under a CES production function, the interpretation of the explosion threshold with automation instead just uses a different primitive: explosive growth requires a sufficient *weight* on AI in its deployment across research and production.

7 Conclusion

> "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)

This paper develops a model for understanding how advances in artificial intelligence could transform economic growth. By considering the interconnected roles of hardware, software, and general technological progress, we show that automation of research and development could generate feedback effects leading to explosive growth. Using historical estimates of diminishing returns, our model suggests that these effects could be substantial, particularly from the automation of computing hardware R&D.

These findings have several implications.

First, 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 central 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).

Second, our analysis suggests that monitoring automation levels in AI R&D 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. This is something economists at AI companies could measure and share publicly.

There is therefore significant scope for further empirical work. These models ultimately require very few parameters, with just a handful pinning down both the explosive growth conditions and the balanced growth paths. Most important is the degree of diminishing returns to research, for which we rely on fairly limited evidence.

References

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."

Epoch AI. 2026. "Trends in Artificial Intelligence." Updated February 5, 2026. Accessed March 29, 2026.

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. 2026. "A.I., Automation, and Explosive Growth: A Condition for Finite-Time Blow-Up." *https://web.stanford.edu/~chadj/explosivegrowthChatGPT5.2.pdf*, Research Note.

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.

Krugman, Paul. 1998. "Two cheers for formalism." *The Economic Journal*, 108(451): 1829–1836.

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.

Roodman, David. 2020. "On the probability distribution of long-term changes in the growth rate of the global economy: An outside view." *Open Philanthropy Report*.

Sandberg, Anders, and David Manheim. 2021. "What Is the Upper Limit of Value?"

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 Proofs

A.1 Deriving conditions for explosive 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', $\boldsymbol{\Omega} := \mathbf{S} - \mathbf{B}$. Therefore, We can relate the (balanced growth path) spillover matrix, $\mathbf{F}$ and the exponent matrix $\boldsymbol{\Omega}$ according to $\boldsymbol{\Omega} = \mathbf{B}(\mathbf{F}-\mathbf{I})$. 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

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

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.

$$\begin{aligned}
\dot{g}_{A,\underline{i}}(t) &= g_{A,\underline{i}}(t)(\boldsymbol{\Omega}g_A + diag(\ell)g_E)_{\underline{i}} && (68)\\
&\geq g_{A,\underline{i}}(t)(\boldsymbol{\Omega}g_A)_{\underline{i}} && (69)\\
&= g_{A,\underline{i}}(t)(\sum_{j\neq\underline{i}} \boldsymbol{\Omega}_{\underline{i},j}g_{A,j} + \boldsymbol{\Omega}_{\underline{i},\underline{i}}g_{A,\underline{i}}(t)) && (70)\\
&\geq g_{A,\underline{i}}(t)(\underline{h}(t)\sum_{j\neq\underline{i}} \boldsymbol{\Omega}_{\underline{i},j}u_j + \boldsymbol{\Omega}_{\underline{i},\underline{i}}g_{A,\underline{i}}(t)) && (71)\\
&= \underline{h}(t)^2 u_{\underline{i}}(\boldsymbol{\Omega}u)_{\underline{i}} && (72)
\end{aligned}$$

where (71) comes from the fact that by definition $g_{A,j}(t) \geq h(t)u_j$ as well as that $\Omega_{\underline{i},j} \geq 0$ and (72) 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

$$\begin{aligned}
\dot{\underline{h}}(t) &\geq \underline{h}(t)^2(\mathbf{\Omega}u)_{\underline{i}} && (73)\\
&= \underline{h}(t)^2(\mu\mathbf{B}u)_{\underline{i}} && (74)\\
&= \underline{h}(t)^2\mu b_{\underline{i}}u_{\underline{i}} && (75)
\end{aligned}$$

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 for all 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

$$\begin{aligned}
\dot{g}_{A,\bar{i}} &= g_{A,\bar{i}}(t)(\mathbf{\Omega}g_A(t) + diag(\ell)g_E)_{\bar{i}} && (76)\\
&= g_{A,\bar{i}}(t)(\mathbf{\Omega}g_A)_{\bar{i}} + g_{A,\bar{i}}(t)\ell_i g_E && (77)\\
&= g_{A,\bar{i}}(t)(\sum_{j\neq\bar{i}}\mathbf{\Omega}_{\bar{i},j}g_{A,j} + \mathbf{\Omega}_{\bar{i},\bar{i}}g_{A,\bar{i}}) + g_{A,\bar{i}}(t)\ell_{\bar{i}}g_E && (78)\\
&\leq g_{A,\bar{i}}(t)(\sum_{j\neq\bar{i}}\mathbf{\Omega}_{\bar{i},j}\bar{h}(t)u_j + \mathbf{\Omega}_{\bar{i},\bar{i}}g_{A,\bar{i}}) + g_{A,\bar{i}}(t)\ell_{\bar{i}}g_E && (79)\\
&= g_{A,\bar{i}}(t)\bar{h}(t)(\sum_{j\neq\bar{i}}\mathbf{\Omega}_{\bar{i},j}u_j + \mathbf{\Omega}_{\bar{i},\bar{i}}u_{\bar{i}}) + \bar{h}(t)u_{\bar{i}}\ell_{\bar{i}}g_E && (80)\\
&= \bar{h}(t)^2u_{\bar{i}}(\mathbf{\Omega}u)_{\bar{i}} + \bar{h}(t)u_{\bar{i}}\ell_{\bar{i}}g_E && (81)\\
&= \bar{h}(t)^2u_{\bar{i}}\mu(\mathbf{B}u)_{\bar{i}} + \bar{h}(t)u_{\bar{i}}\ell_{\bar{i}}g_E && (82)
\end{aligned}$$

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

$$\dot{\bar{h}}(t) \leq \bar{h}(t)^2\mu b_{\bar{i}}u_{\bar{i}} + \bar{h}(t)\ell_{\bar{i}}g_E \quad (83)$$

and since $\mu < 0$ the quadratic part of the expression dominates as $\bar{h}$ grows so $\bar{h}(t)$ 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 $g_A^{\text{BGP}} \propto u$. □

*Proof of Proposition 1.**. First, note that $\mathbf{F}^Y$ and $\mathbf{F}^A$ terms (and therefore $\mathbf{\Psi}_{A,Y}$) can be derived from combining equations (30) and (31).

Next, we focus on demonstrating that Proposition 1 can be directly applied to a system with capital, in particular, to show that depreciation does not affect whether the system is explosive. We begin by following the method suggested in Jones (2026). Specifically, define

$$a_i(t) = A_i(t)^{\beta_i} \qquad \text{and} \qquad k(t) = K(t)^{1-\alpha}.$$

Taking time derivatives and using the definitions of $\dot{K}$ and $\dot{A}_i$ from (24) and (26) yields

$$\begin{aligned} \dot{k}(t) &= (1-\alpha)K(t)^{-\alpha}\dot{K}(t) && (84) \\ &= (1-\alpha)a\bar{A}s_Y L_Y^{1-\alpha} - (1-\alpha)\delta K(t)^{1-\alpha} && (85) \\ &= c_1 \prod_{i \in I} a_i(t)^{\frac{\tau_i}{\beta_i}} - \delta_1 k(t), && (86) \end{aligned}$$

where $c_1 = (1-\alpha)as_Y L_Y^{1-\alpha}$, and $\delta_1 = (1-\alpha)\delta$. Similarly,

$$\begin{aligned} \dot{a}_i(t) &= \beta_i A_i(t)^{\beta_i - 1}\dot{A}_i(t) && (87) \\ &= c_{2,i} k(t)^{\frac{\gamma_i \lambda_i}{1-\alpha}} \prod_{j \in I \setminus i} a_j(t)^{\frac{\phi_{i,j}}{\beta_j}}, && (88) \end{aligned}$$

where $c_{2,i} = \beta_i v_i \kappa_i^{\lambda_i \gamma_i} L_i^{\gamma_i \lambda_i}$ and κ_i is the share of capital allocated to i.

Next, define $\tilde{k}(t) = e^{\delta_1 t}k(t)$, so $\dot{\tilde{k}}(t) = [\dot{k}(t) + \delta_1 k(t)]e^{\delta_1 t}$. Rearranging (86) gives

$$\begin{aligned} c_1 \prod_{i \in I} a_i(t)^{\frac{\tau_i}{\beta_i}} &= \dot{k}(t) + \delta_1 k(t) && (89) \\ &= \dot{\tilde{k}}(t)e^{-\delta_1 t}, && (90) \end{aligned}$$

and substituting $k(t) = \tilde{k}(t)e^{-\delta_1 t}$ into (88) gives

$$\dot{a}_i(t) = c_{2,i}\left(\tilde{k}(t)e^{-\delta_1 t}\right)^{\frac{\gamma_i \lambda_i}{1-\alpha}} \prod_{j \in I \setminus i} a_j(t)^{\frac{\phi_{i,j}}{\beta_j}}. \qquad (91)$$

Therefore, we obtain the system

$$\dot{\tilde{k}}(t) = e^{\delta_1 t} c_1 \prod_{i \in I} a_i(t)^{\frac{\tau_i}{\beta_i}}, \tag{92}$$

$$\dot{a}_i(t) = c_{2,i} \left(\tilde{k}(t) e^{-\delta_1 t}\right)^{\frac{\gamma_i \lambda_i}{1-\alpha}} \prod_{j \in I \setminus i} a_j(t)^{\frac{\phi_{i,j}}{\beta_j}}. \tag{93}$$

We observe that (92)–(93) is almost of the form required for Proposition 1, except for the time-varying exponential terms. To make the terms match exactly, define

$$b_i := \frac{\gamma_i \lambda_i}{1-\alpha} > 0 .$$

Now define an new system $(\tilde{\tilde{k}}, \tilde{a}; \nu, \Gamma)$ where $\nu = (\nu_1, \{\nu_{2,i}\}_{i \in I})$ is a collection of positive scalars and Γ collects the structural parameters (which determine the exponents):

$$\dot{\tilde{\tilde{k}}}(t) = \nu_1 c_1 \prod_{i \in I} \tilde{a}_i(t)^{\frac{\tau_i}{\beta_i}}, \tag{94}$$

$$\dot{\tilde{a}}_i(t) = \nu_{2,i} c_{2,i} \, \tilde{\tilde{k}}(t)^{b_i} \prod_{j \in I \setminus i} \tilde{a}_j(t)^{\frac{\phi_{i,j}}{\beta_j}}. \tag{95}$$

Proposition 1 applies directly to (94)–(95) (after treating the capital as a new dimension of the technology system), and whether this system is explosive depends on the exponent structure (i.e. on Γ through the exponent matrix), and not on the particular positive coefficients ν.

It remains to show that (92)–(93) is explosive if and only if (94)–(95) is explosive. Fix any horizon $T > 0$. Over $t \in [0, T]$ we have the uniform bounds

$$1 \le e^{\delta_1 t} \le e^{\delta_1 T}, \qquad e^{-\delta_1 T b_i} \le e^{-\delta_1 t b_i} \le 1.$$

Define two upper $(+)$ and lower $(-)$ bounding systems with the same exponents as

(92)–(93):

$$\dot{\tilde{\tilde{k}}}^{-}(t) = 1 \cdot c_1 \prod_{i \in I} (\tilde{a}_i^{-}(t))^{\frac{\tau_i}{\beta_i}}, \tag{96}$$

$$\dot{\tilde{a}}_i^{-}(t) = e^{-\delta_1 T b_i} \cdot c_{2,i} \, (\tilde{\tilde{k}}^{-}(t))^{b_i} \prod_{j \in I \setminus i} (\tilde{a}_j^{-}(t))^{\frac{\phi_{i,j}}{\beta_j}}, \tag{97}$$

$$\dot{\tilde{\tilde{k}}}^{+}(t) = e^{\delta_1 T} \cdot c_1 \prod_{i \in I} (\tilde{a}_i^{+}(t))^{\frac{\tau_i}{\beta_i}}, \tag{98}$$

$$\dot{\tilde{a}}_i^{+}(t) = 1 \cdot c_{2,i} \, (\tilde{\tilde{k}}^{+}(t))^{b_i} \prod_{j \in I \setminus i} (\tilde{a}_j^{+}(t))^{\frac{\phi_{i,j}}{\beta_j}}. \tag{99}$$

Given the definitions then of these systems we have

$$(\tilde{\tilde{k}}^{-}(t), \tilde{a}^{-}(t)) \leq (\tilde{k}(t), a(t)) \leq (\tilde{\tilde{k}}^{+}(t), \tilde{a}^{+}(t)) \quad \text{for all } t \in [0, T].$$

In particular, if the lower system is explosive by time T, then so is (92)–(93); and if (92)–(93) is explosive by time T, then so is the upper system.

By Proposition 1, explosiveness for (94)–(95) depends on the exponent structure in Γ and not on the choice of positive coefficients ν. Therefore the lower and upper comparison systems are explosive if and only if $(\tilde{\tilde{k}}, \tilde{a}; \nu, \Gamma)$ is explosive for any $\nu > 0$. Combining this with the sandwiching argument above implies that (92)–(93) is explosive if and only if $(\tilde{\tilde{k}}, \tilde{a}; \nu, \Gamma)$ is explosive. Since δ enters (92)–(93) only through the bounded positive factors $e^{\pm \delta_1 t}$, it follows that depreciation does not affect whether the system is explosive.

Then, since $\tilde{k}(t) = e^{\delta_1 t} k(t)$ and $k(t) = K(t)^{1-\alpha}$ are strictly increasing transformations, and similarly $a_i(t) = A_i(t)^{\beta_i}$ is strictly increasing in $A_i(t)$, the system in (35)–(42) is explosive if and only if (92)–(93) is explosive.

Therefore, we have proved that one can apply Proposition 1 to consider the dynamics of (23) - (28). However, in doing so, the above proof requires treating capital as an additional technology, and considering the eigenvalues of an $(N+1) \times (N+1)$ matrix:

$$\tilde{\mathbf{F}} = \begin{pmatrix} 0 & p' \\ q & \mathbf{G} \end{pmatrix}, \tag{100}$$

where $p_i = \tau_i / \beta_i$, $q_i = \gamma_i \lambda_i / (1 - \alpha)$, and $\mathbf{G}_{i,j} = \phi_{i,j} / \beta_j$. We proceed to show we can instead look at eigenvalues of $\mathbf{F}^A + \mathbf{F}^Y$. Taking the definitions of $\mathbf{F}^A$ and $\mathbf{F}^Y$ we have

$$\mathbf{G} = \mathbf{B}\mathbf{F}^A\mathbf{B}^{-1} \quad \text{and} \quad dc' = \mathbf{B}\mathbf{F}^Y\mathbf{B}^{-1}$$

where $\mathbf{B} = diag(\beta_1, \ldots, \beta_N)$. Therefore,

$$\mathbf{G} + qp' = \mathbf{B}(\mathbf{F}^A + \mathbf{F}^Y)\mathbf{B}^{-1}$$

and since these two matrix sums are similar then $\rho(\mathbf{G} + qp') = \rho(\mathbf{F}^A + \mathbf{F}^Y)$. Therefore, we just have to relate the eigenvalues of $\tilde{\mathbf{F}}$ and $\mathbf{G} + qp'$. To this end, note, the usual eigenvalue and eigenvector ($\mathbf{x} = (x_0, \bar{\mathbf{x}})'$, where x_0 is scalar and $\bar{\mathbf{x}}$ is an N length vector) decomposition:

$$\begin{aligned} \lambda\mathbf{x} &= \tilde{\mathbf{F}}\mathbf{x} & (101) \\ &= \begin{pmatrix} p'\bar{\mathbf{x}} \\ \\ qx_0 + \mathbf{G}\bar{\mathbf{x}} \end{pmatrix} & (102) \end{aligned}$$

and solving the top row we have $x_0 = \lambda^{-1}p'\bar{\mathbf{x}}$. Therefore we have the dimension reduced decomposition

$$(\mathbf{G} + \frac{1}{\lambda}qp')\bar{\mathbf{x}} = \lambda\bar{\mathbf{x}} \tag{103}$$

Note, that by Perron-Fronbius, since dc' is non-negative then $\rho(\mathbf{G} + \frac{1}{\lambda}qp')$ is decreasing in λ. Therefore

$$\begin{cases} \rho(\mathbf{G} + \frac{1}{\lambda}qp') > \rho(\mathbf{G} + qp') & \text{if} \quad \lambda < 1 \\ \rho(\mathbf{G} + \frac{1}{\lambda}qp') = \rho(\mathbf{G} + qp') & \text{if} \quad \lambda = 1 \\ \rho(\mathbf{G} + \frac{1}{\lambda}qp') < \rho(\mathbf{G} + qp') & \text{if} \quad \lambda > 1 \end{cases} \tag{104}$$

and from this we can combine with above results for the chain

$$\rho(\tilde{\mathbf{F}}) \gtrless 1 \iff \rho(\mathbf{G} + qp') \gtrless 1 \iff \rho(\mathbf{F}^A + \mathbf{F}^Y) \gtrless 1 \; .$$

□

A.2 Applying proposition 1 to corollaries 1 and 2

Here we the proof for corollary 2. Corollary 1 is a special case of corollary 2 setting $f_Y = 0$ and $\tau = 0$.

Proof of Corollary 2. From above, we have

$$\hat{\mathbf{F}}^A = u\sigma' - D \tag{105}$$

$$\hat{\mathbf{F}}^Y = uv' \tag{106}$$

with

$$u_i := \frac{f_i \lambda_i}{\beta_i - f_i \lambda_i \sigma_i}, \qquad v := \frac{1}{1-f_Y}\Big(\frac{\tau}{1-\alpha} + f_Y \sigma\Big) \qquad D = diag(u \odot \sigma)$$

and $\odot$ is an elementwise multiplication. Then, we have

$$\hat{\mathbf{F}} := \hat{\mathbf{F}}^A + \hat{\mathbf{F}}^Y = uw' - D \tag{107}$$

where $w = \sigma + v$ and σ is the vector of σ_i terms.

Next, the standard matrix determinant lemma yields

$$\det\Big(\lambda I - \hat{\mathbf{F}}\Big) = \det(\lambda I + D)\Big(1 - w^\top (\lambda I + D)^{-1} u\Big).$$

and for λ to be an eigenvalue of $\hat{\mathbf{F}}$ we require the RHS to equal zero. Since D is a diagonal matrix with strictly positive entries, $\det(\lambda I + D) > 0$. Therefore we require

$$1 = w^\top (\lambda I + D)^{-1} u \tag{108}$$

$$= \sum_{i \in I} \frac{w_i u_i}{\lambda + u_i \sigma_i} \tag{109}$$

Notice, for $\lambda > 0$ the RHS of this equations is strictly decreasing in λ, hence there is at most one $\lambda > 0$ to satisfy this equality. Therefore, $\hat{\mathbf{F}}$ has at most one positive eigenvalue, and by Perron-Frobenius $\rho(\hat{\mathbf{F}}) = \lambda$. To conclude the proof, if $\lambda > 0$ exists then,

$$\rho(\hat{\mathbf{F}}) > 1 \iff \lambda > 1 \iff \sum_{i \in I} \frac{w_i u_i}{1 + u_i \sigma_i} > 1 \iff f_Y + \sum_{i \in I} f_i r_i \Big(\frac{\tau_i}{1-\alpha} + \sigma_i\Big) > 1$$

where the middle iff comes from (109) being equal to one and strictly decreasing in λ, so if $\lambda > 1$ substituting in 1 for λ increases (109). The last iff comes from substituting out w and u from summation term. Thus, Proposition 1 directly applies. □

B Balanced and hyperbolic growth with a fixed factor

Here we take the growth model from section 4.4, 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

$$Y = \bar{A} K_Y^{\alpha_K} \hat{L}_Y^{\alpha_L} M^{1-\alpha_K-\alpha_L}. \tag{110}$$

We maintain equations (37) - (42) 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

$$g_A^{\mathrm{BGP}} = \hat{\mathbf{\Psi}}_{A,Y}\left(\hat{\mathbf{r}}\frac{1}{1-\alpha_K-\alpha_L f_Y}(\mathbf{1}_n - f - \alpha_K(\mathbf{1}_n - f) - \alpha_L(\mathbf{1}_n \times f_Y - f))\right)\mathbf{n} \tag{111}$$

where

$$\hat{\mathbf{\Psi}}_{A,Y} = (\mathbf{I} - [\hat{\mathbf{F}}^A + \hat{\mathbf{F}}^Y])^{-1} \tag{112}$$

$$\hat{\mathbf{F}}^Y := \underbrace{[\hat{\mathbf{r}} f]}_{dg_A^{\mathrm{BGP}}/dg_Y^{\mathrm{BGP}} \text{ from tech L.O.M}} \underbrace{\left[\frac{\alpha_L}{1-\alpha_K-\alpha_L f_Y} \times [\frac{1}{\alpha_L}\tau + f_Y\sigma]\right]'}_{dg_Y^{\mathrm{BGP}}/dg_A^{\mathrm{BGP}} \text{ from } 110} \tag{113}$$

and $\hat{\mathbf{r}}$ and $\hat{\mathbf{F}}^A$ are defined as in section 4.3.

Then, deriving the explosive growth condition, we get

Corollary 5. *The automation-calibrated growth model with a fixed factor (described in equations (110) and (37) - (42)) explodes in finite time iff*

$$\frac{\alpha_L}{1-\alpha_K} f_Y + \sum_{i\in I} f_i r_i \left(\frac{\tau_i}{1-\alpha_K} + \sigma_i\right) > 1 \tag{114}$$

C Additional figures

Figure 11: Automation Shock Experiment (Low BGP: $n_i = 1\%$ for all i)

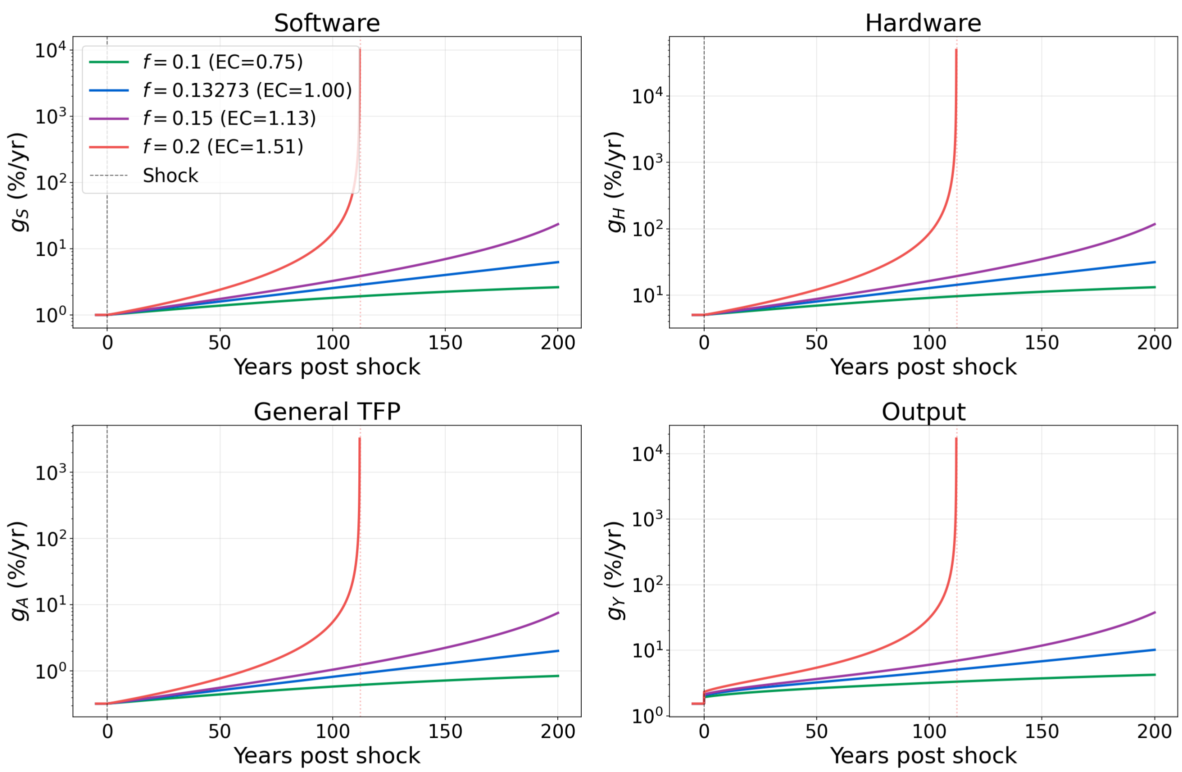

Note: 'EC' reports the values of $f(1+r_S+r_H+\frac{r_A}{1-\alpha})$, where r and α values are taken from table 3. Here the balanced growth path conditions imply savings rates $(a_Y, a_H) = (0.34, 0.11)$.