(Artificial) Intelligence Saturation and the Future of Work

[PRELIMINARY; DO NOT CIRCULATE]

Konrad Kording and Ioana Marinescu

August 1, 2025

Abstract

Predicting how the advent of artificial intelligence (AI) will affect production and the labor market is key for policy design. We build a tractable production function framework that separates *physical* and *intelligence* inputs and allows the marginal returns to intelligence to *saturate*. Whereas physical capabilities let us do things in the world, intelligence capabilities let us do them well; we assume that the latter can display sharply diminishing returns. Because the price of AI capital is falling much faster than that of physical capital, intelligence tasks are automated first, pushing human labor toward physical input production and niche intellectual tasks. We present an aggregate CES-type production function that nests a standard capital–labor model when AI is inexistent, and exhibits economy-wide intelligence saturation when AI is abundant. The resulting dynamics are often non-monotonic. In the baseline calibration, wages increase and then decrease with automation. Wage decreases are steeper when physical and intelligence inputs are more substitutable. After full automation, more AI but above all more physical capital increase wages, a classic prediction from standard production functions in capital and labor. Yet, the limits imposed by intelligence saturation blunt the long-run wage trajectory, while returns to scarce physical capital keep climbing. Policy therefore faces a dual challenge: (i) pace automation to blunt a potential short-run wage decline and (ii) accelerate labor-augmenting investment in the physical sector to lift the post-automation wage floor.

1 Introduction

OpenAI's public release of ChatGPT (Nov 30 2022) demonstrated to the public the power of new AI technologies Brown et al. (2020), and the advent of capable large language models ushered a new era of uncertainty and debate about the future of work. This debate is dominated by two groups with rather different ways of thinking, AI experts and economists. This paper is co-written by one author from each field and aims at providing a high-level consensus.

On the one hand, many tech experts boldly predict that most human labor can be replaced by AI, ushering an astounding new era of economic growth – sometimes dubbed

"singularity" Kurzweil (2005). This stance is probably culturally driven in part by the fact that many indicators in AI are exponential with very short doubling times, often of the order of months. If computers can become twice as fast every two years and tokens for large language models can become half the price within a few months, why shouldn't the economy undergo similar trajectories?

On the other hand, many economists cautiously foresee a more vanilla future: AI is a new general purpose technology, with positive but ultimately limited effects on growth, and with some negative but ultimately transitory effects on the labor force. Economists arrived at this stance observing a slow but steady growth since the dawn of the Industrial Revolution, partially driven by an accumulation of general purpose technologies such as electricity, and more generally an innovation economy.

Which view is closer to the truth is highly consequential for workers' welfare and for policy. If the full labor replacement view is closer to the truth, sweeping new policies may be needed to insure that displaced workers can continue to receive an income. If the new but limited general purpose technology view is correct, there is less concern about the future of work, while policies to smooth the adaptation to the new technology may still be in order. We tackle the overarching question: When does automating cognition help workers, and when does it harm them? And what are the relevant timescales?

Here we aim to use economic theory to bridge the two views and illustrate plausible scenarios. Following the takes of the AI folks, we allow AI to replace humans in many tasks: specifically, we assume AI may fully replace humans in intelligence tasks, i.e. those that can be done in a disembodied way, or virtually. These disembodied intelligence tasks plausibly represent a large share of all tasks, since 60% of workers could telework at the height of the COVID-19 pandemic (Barrero et al., 2021). We even acknowledge that there are domains where the production of artificial intelligence capital may grow at rates typical in computer science but unheard of in economics. Following the takes of the economics folks, we argue that the returns to intelligence saturate, which limits the growth and wage effects of AI. As an example of the saturating impact of intelligence, consider the decreasing impact of scientific research on economic growth (Bloom et al., 2020). The low-hanging fruit has likely been harvested, so research and patents are getting less disruptive over time (Park et al., 2023), with progress especially difficult in fields where knowledge is already deep (Jones, 2009). It is notable though that patents in computers and electronics show increasing innovativeness since the 1980s (Kelly et al., 2021), which contributes to the expansion of AI capabilities.

The key focus of our paper is on intelligence saturation, the idea that you build things through physical means, and intelligence can at best make you maximally efficient at that. Both physical and intelligence components are needed in production. To illustrate, take education: it has both a physical and an intellectual component. Quality education requires a physical component, the teacher and the students need to be in the same room for efficient learning. Consequently, COVID-19 lockdowns negatively impacted student academic performance Cortés-Albornoz et al. (2023), and districts with full in person learning fared significantly better Jack et al. (2023). Therefore, we need teachers and classrooms – the physical component. Of course, education also has an intellectual component: the teaching materials. AI could improve teaching materials, but it is likely that such improvements would saturate as the teacher ultimately needs to deliver the materials effectively, and the classroom setup needs to be conducive to learning. Similarly, the production of cars needs

machines and controllers, the production of health requires physical components and intellectual components, and the production of food requires a cook and a recipe. No amount of intelligence in controllers can dramatically speed up a modern production line. No amount of intelligence can make doctors much more efficient and as far as we know there are no recipes that dramatically improve the enjoyment we get out of food. We see intelligence saturation as unfolding across the many tasks in the economy.

To capture these insights, our theoretical approach uses an aggregate production function (how the output of the economy depends on inputs like labor) that is a Constant Elasticity of Substitution (CES) function of physical and intelligence inputs. The output is driven by a combination of Physical and Intellectual inputs. The Physical and Intellectual inputs now consist of Physical and Intellectual capital, both of which are augmented by human labor which is assumed to flow freely between the two domains. The resulting model, a nested CES type model, allows us to then quantify how changing AI inputs, both in terms of which tasks can be automated and how much AI capital exists over time, affects production and wages.

We make the standard assumption about capital-labor substitution. In intelligence tasks, AI and labor are good substitutes, like in the task model Acemoglu and Restrepo (2018); Acemoglu et al. (2024) also used by Jones and Liu (2024); Korinek and Suh (2024). The difference is that we only assume good substitution for intelligence tasks, not for all tasks. We assume that substitution is considerably harder in the physical domain. The scaling of capabilities of robots with their limited bodies is simply much slower than the scaling of pure intelligence. This asymmetric treatment of capital-labor in intelligence vs physical tasks mirrors the long-run differences in learning curves in the two areas. Our model is a relatively standard model but for the fact that we fundamentally enshrine the complementarity of physical and intellectual inputs.

We calibrate our model using parameters estimated in the economics literature and find that automation in the intelligence block first increases and then decreases wages: at first, productivity gains in intelligence more than offset the wage pressure from workers crowding into fewer tasks, but eventually the negative effects dominate as most workers get iced out of intelligence tasks. This result presents a cautionary tale: initial wage increases could turn into wage declines as automation progresses.

We then analyze the dependency of the model on parameter changes focusing on effects of automation on wage and production. Stronger decreasing returns in intelligence dampen wage effects. Lower substitutability between intelligence tasks leads to a more pronounced hump-shaped pattern in wages. The interactions between the physical and intelligence sectors are critical for the impact of automation on wages and output because workers crowd in the physical sector as automation progresses. Higher substitutability between the physical and intelligence blocks makes the wage effects of automation more negative but the output effects more positive. A lower share for the physical sector leads to more negative impacts of automation on wages, while output benefits. Intuitively, it is bad for wages if physical goods and in person services can easily be substituted with fully disembodied or virtual intelligence products. And having a lot of the economy in the physical sector means workers have a bigger place to land as automation progresses. But from the output perspective, the conclusions are opposite: for the economy to maximally benefit from automation and AI, we need more substitution between physical and intelligence, and a lower share for the physical

sector. This situation would allow the economy to be dominated by the intelligence sector where AI can contribute the most to productivity.

From a policy perspective, our model suggests that reducing any negative effect of automation on wages could be accomplished by slowing down automation while stimulating investment in physical capital. Slowing down automation obviously reduces wage losses but is costly in terms of foregone growth (but less costly the more intelligence saturation there is), and it could be seen as gaining time to make necessary investments in physical capital. Physical capital investments are helpful because they increase the marginal productivity of labor in the physical task, and therefore prop up wages as workers transition into the physical sector. An additional policy lever would be the taxing of virtual substitutes to in-person services or physical goods. Such a policy would prevent increasing substitutability between physical and intellectual inputs, which would increase the extent to which workers in the physical sector benefit from automation and from further increases in AI capital.

We make three contributions to the literature. First, we add to the macro growth and production function literature (Ray et al., 2022) by explicitly considering the role of intelligence as an input into production rather than an overall technology factor. Modeling intelligence as a separate factor with saturating returns reveals an endogenous upper bound on AI-driven growth. Second, we embed the task-based view of skill-biased technical change Acemoglu and Autor (2011) in a macro setting where automation acts asymmetrically across physical and intelligence tasks, linking micro substitution patterns to macro outcomes. Third, we provide a unified lens on the future-of-work debate by delivering clear, testable conditions – rooted in substitution elasticities and labor shares – under which wages increase or decrease as automation progresses. Collectively, these contributions clarify why AI can be transformative without boundless growth and how policy can shape the distribution of its gains. Put simply, we ask: When does substituting human cognition with machines help workers, and when does it hurt them?

To help readers intuit the interacting components of our model we make available a web-app that makes it easy to analyze the effect of the various parameters.

In section 2, we discuss the background behind the assumption of intelligence saturation. In section 3, we discuss the modeling setup, and the theoretical impacts of automation and AI on wages (in section 3.3). We provide model simulations in section 4. Our discussion is in section 5. Finally, section 6 concludes.

2 Why intelligence saturation

2.1 Scaling of AI, physical capital, and labor and why physical and intellectual inputs should be modeled differently

The past division of the economy into labor and capital inputs has served the economics discipline well for decades, why do we need to divide it into physical and intellectual inputs? Here we will make the argument that physical and intellectual tasks at this point of time in history scale in fundamentally different ways, and that therefore this distinction is necessary to be able to drive meaningful predictions about AI.

Computer science has long scaled in a way that is entirely different from the scaling of

the rest of the economy. Moore's law states that the number of transistors on a microprocessor doubles every 2 years. This scaling has been going on for more than 70 years. And while it is ending, microprocessors are now rapidly getting better through other channels, e.g. multicores. Computation thus becomes cheaper, roughly by a factor of two every two yearsNordhaus (2007). This background with scaling gave the computer science a focus on how scaling affects systems.

This scaling is faster in AI than in computer science. At the moment, for a given task that AI can solve, e.g. text generation, the price decreases rapidly. The price currently halves roughly every 3 months DeepLearning.AI (2024). These decreases come from multiple sources. First, compute becomes cheaper by a factor of two roughly every two years since the 1950s Nordhaus (2007). Second, compute becomes cheaper in terms of energy, with energy efficiency improving roughly two orders of magnitude per decade Koomey et al. (2011). Third, improving AI algorithms makes running AI cheaper Hernandez and Brown (2020). There are no indications that this trend is changing Guo et al. (2025). Clearly the price of AI is decreasing exponentially at a rate much faster than pretty much anything else we know in the economy. This makes it more and more economically advantageous to use AI for tasks where it can be used.

We can also talk about the scaling of labor and physical capital. Labor in the US is increased through population growth (currently less than 1% per year) and supply of higher quality labor (but education gains are slowing down). Physical capital increases through our economy's savings rate and its growth is currently estimated to be in the 3% range. Relatively to the breakneck speed of AI developments, the other economic inputs are basically stationary.

A special case of the scaling of capital is the scaling of robotics; progress at this endeavor could make human physical labor far more substitutable. However, the price of robots, like that of most goods, tends to rather decrease as a power law than exponential. The average price of industrial robots decreased from about $ 100k to about $ 20k from 1995 to 2025 and the decrease is slowing down considerably ARK Investment Management LLC (2021). Generally, when it comes to physical goods, prices tend to transition from an early phase of rapid price decline towards a very steady long term price (e.g. U.S. Bureau of Labor Statistics (2025b)). Outside of an early decrease in price, physical good prices tend to stagnate over time.

Now, if computer tasks become cheap exponentially at a very high rate, why did this so far not carry through into the economy, giving much higher growth rates? The reason was two fold: (1) computer services are traditionally a relatively small part of the economy (2) their role in the production function was limited to a certain number of tasks, without much substitutability in most tasks. Both factors are changing with the AI share in the economy increasing and substitutability getting much higher because AI can progressively solve more complex intllectual tasks. This observation is what in our opinion necessitates a shift towards separating the modeling of the physical (where growth should arguably look somewhat similar to the past) from the modeling of the intellectual (where, arguably, growth may be exceptionally fast). This paper is our stab at modeling these two contrasting realms.

2.2 Intelligence saturation: Modeling production as a function of physical and intelligence inputs

Traditionally, economists discuss the production function in terms of capital and labor. This view acknowledges that production of goods usually has a component that requires capital, say trucks, or desks, and a component that requires labor, say driving those trucks or solving computer science problems on those desks. Capital is usually abbreviated as K while labor is abbreviated as L. Popular conceptualizations of their interactions then gives rise to production functions for tasks: the amount of goods produced in a task i is a function of the amounts of capital K_i and labor L_i allocated to the it. The micro or local production function for tasks is then $Y_i = f(K_i, L_i)$. Economists usually assume that capital and labor are not perfectly substitutable, e.g. through a CES function (like $\left[\alpha\, K_i^{\rho} \;+\; (1-\alpha)\, L_i^{\rho}\right]^{1/\rho}$). The intuition behind this function is that adding more capital or more labor both increase output, but the two are not perfectly substitutable, so that the marginal productivity of either decreases when the other is held fixed; the elasticity of substitution is regulated by ρ.

Besides the different pricing trends in AI vs. physical capital discussed above, we see fundamental differences between physical and intelligence inputs based on models of control in robotics, psychology, and neuroscience. Within these areas, we talk about effectors. An effector can be thought of as something physical that has an effect on the world. An arm is an effector. As is the shovel of an excavator. Effectors have the capability to change the world. More and bigger effectors have the capability of changing the world more. And for every effector and task, we use intelligence to make these changes be as positive for us as possible, and that is what AI is about. More intelligence allows the effectors to be less wasteful. In the limit of perfect intelligence, every use of an effector is as efficient as it can be, there is no more waste. In this AI framework intelligence saturation is natural.

Here we thus change the division of the economy into capital and labor, adding an extra dimension of physical vs intelligence inputs. Physical tasks can be done by machines, requiring capital, or by humans, requiring labor. And intelligence tasks can be done by machines in the form of AI, requiring capital for compute and investment into enabling this possibility, and they can also be done by human labor. In this paper, we will argue that this division of the economy into four groups (human vs machine, intelligence vs physical) allows us to more precisely model the economy as it exists and to better predict what is to come.

3 Models

3.1 Marginal productivity of intelligence, AI, and labor

The marginal productivity of K_I is:

$$\frac{dY}{dK_I} = \frac{\partial F}{\partial I}\frac{dI}{dK_I} \tag{1}$$

This equation implies that there are two ways to model a decreasing marginal productivity of AI: either reduce the output gains of adding more intelligence ($\frac{\partial F}{\partial I}$) or reduce the effectiveness of AI in producing more intelligence ($\frac{dI}{dK_I}$). We will explore both of these avenues.

As long as we haven't automated all intellectual tasks, labor is used in the production of both physical and intelligence inputs, so the marginal productivity of labor is:

$$\frac{dY}{dL} = \frac{\partial F}{\partial P}\frac{dP}{dL} + \frac{\partial F}{\partial I}\frac{dI}{dL} \tag{2}$$

Once full automation of intellectual tasks is achieved, labor is used only in the production of physical inputs, so the marginal productivity of labor is:

$$\frac{dY}{dL} = \frac{\partial F}{\partial P}\frac{dP}{dL} \tag{3}$$

At face value, dropping the second term $\frac{\partial F}{\partial I}\frac{dI}{dL}$ means the wage must fall, and this is indeed the intuition most people have about automation. A rapid shift to full automation can even magnify the drop, because with fixed physical capital the physical-block marginal product falls as labor floods in.

However, the first and the second equation above are evaluated at different input mixes P, I, so they are not directly comparable. If AI is so abundant that automating the very last tasks still generates large gains in the intelligence input (and if those gains feed strongly into final output), then the marginal productivity in the physical channel $\partial F/\partial P$ can rise enough to more than offset the lost intelligence term. In other terms, if getting to full automation greatly increases I, and this increase sufficiently benefits the productivity of the physical input (where labor resides), you could actually see the full-automation wage end up higher than the no-automation wage.

After full automation is achieved, increasing the amount of AI in the economy unambiguously benefits workers, because more AI increases the physical input productivity (where labor is employed), that is:

$$\frac{\partial}{\partial K_I}\left[\frac{dY}{dL}\right] = \frac{\partial^2 F}{\partial P\,\partial I}\;\frac{dI}{dK_I}\;\frac{dP}{dL} \;>\; 0 \tag{4}$$

The cross-partial $\frac{\partial^2 F}{\partial P\,\partial I}$ is positive for a CES with exponent $\rho < 1$, which is our maintained assumption.

Here it's worth noting that the marginal productivity of labor increases with AI only as long as intelligence increases with AI ($\frac{dI}{dK_I} > 0$). If we have intelligence saturation ($\frac{dI}{dK_I} = 0$), then wages no longer increase with AI.

3.2 CES setup

We assume a *separate* production function for each of the physical and intelligence inputs, each produced with labor and capital. Each of the two blocks is CES-like, and joined by a CES aggregator at the macro level. This lets us vary substitution elasticities and scale parameters independently for physical and intelligence inputs.

Table 1 summarizes the parameters used in this model, together with their domain assumptions. Assumptions on parameter domains are generally standard. To note, intelligence tasks are assumed to be substitutes with each other, i.e. $\rho_I > 0$. This assumption is empirically plausible as we further discuss when setting up our calibration[1] below, and analytically

[1]The model can be solved numerically regardless of the sign of ρ_I.

convenient because some derivations take a different form and the signs of some expressions can switch if $\rho_I < 0$.

Parameter	Notation	Description	Domain / Assumption
Physical capital	K	Aggregate physical capital	$K \geq 0$
AI input	$\mathrm{K_I}$	AI capital used in intel CES	$\mathrm{K_I} \geq 0$
Labor	L	Aggregate labor	$L \geq 0$
Labor share to physical tasks	α	$L_P = \alpha L$	$0 \leq \alpha \leq 1$
Labor to physical block	L_P	Labor in physical: $L_P = \alpha L$	$L_P \geq 0$
Labor to intel block	L_I	$L_I = (1-\alpha)L$	$L_I \geq 0$
Physical CES share	α_P	Weight on K vs. L_P in P	$0 \leq \alpha_P \leq 1$
Physical CES exponent	ρ_P	Exponent in $P = \left(\alpha_P K^{\rho_P} + (1-\alpha_P) L_P^{\rho_P}\right)^{1/\rho_P}$	$\rho_P < 1, \rho_P \neq 0$
Automation fraction	α_I	Fraction of tasks automated	$0 \leq \alpha_I \leq 1$
Intel CES micro-exponent	ρ_I	Exponent in I	$0 < \rho_I < 1$
Intel returns to scale	θ_I	Returns to scale in intel block	$0 < \theta_I \leq 1$
Macro CES share	τ	Weight on P vs. I in final CES	$0 < \tau < 1$
Macro CES exponent	ρ	Exponent in $Y = \left[\tau P^{\rho} + (1-\tau) I^{\rho}\right]^{1/\rho}$	$\rho < 1, \rho \neq 0$

Table 1: Macro-blocks production parameters

Physical Tasks (P) The production function for physical tasks is given by a a CES in physical capital K and labor L (with only a fraction $\alpha \in [0,1]$ of total labor allocated to P). We may denote $L_P \equiv \alpha\, L$ (and correspondingly $L_I \equiv (1-\alpha)L$ for intellectual tasks).

$$P \;=\; \Big(\alpha_P\, K^{\rho_P} \;+\; (1-\alpha_P)\, L_P^{\rho_P}\Big)^{1/\rho_P}. \tag{5}$$

Intellectual Tasks (I) For the intelligence input, we use the type of production function that Jones and Liu (2024); Korinek and Suh (2024) use for the whole economy. It is a CES in intelligence tasks (e.g. write text, make tables, etc.).

$$I = \left[\int_i y(i)^{\rho_I} d\Phi(i)\right]^{\theta_I/\rho_I}, \quad \rho_I = \frac{\sigma_I - 1}{\sigma_I}, \tag{6}$$

where $\Phi(i)$ reflects the cumulative mass (or fraction) of tasks with a complexity level less than or equal to i, and σ_I is the elasticity of substitution between tasks. θ_I is a returns to scale parameter. In Korinek and Suh (2024), there are constant returns to scale, but we allow for decreasing returns to scale if $\theta_I < 1$.

The task-level production function for intelligence is consistent with the assumptions in Korinek and Suh (2024) and Jones and Liu (2024).

Let $\Phi(i) = \alpha_I \in [0, 1]$ represent the fraction of intellectual tasks that is automated. As shown in Korinek and Suh (2024), the above assumptions lead to a CES function in the inputs:

$$I \;=\; \left[{\alpha_I}^{1-\rho_I} \mathrm{K_I}^{\rho_I} \;+\; (1-\alpha_I)^{1-\rho_I} {L_I}^{\rho_I}\right]^{\theta_I/\rho_I} \tag{7}$$

where $L_I = (1-\alpha)L$ is the labor allocated to intellectual tasks, and $\rho_I > 0$ indicates high substitutability between intelligence tasks. We allow for $\theta_I < 1$, capturing decreasing returns to scale in the production of the intelligence input.

The above equation only obtains if AI is sufficiently abundant so that the wage – aka the marginal productivity of labor – is greater than the marginal productivity of AI. As long as this holds, labor is only used in the not-yet-automated tasks. The condition is:

$$\frac{K_I}{L_I} > \frac{\alpha_I}{1-\alpha_I}\,. \tag{8}$$

We assume that this condition holds: AI is abundant.

Aggregate CES production function with separate physical and intellectual production functions The final output is obtained by aggregating the two intermediate inputs – physical and intellectual – via a CES aggregator:

$$Y = F(P, I) \;=\; \left[\tau\, P^{\rho} \;+\; (1-\tau)\, I^{\rho}\right]^{1/\rho}, \qquad 0 < \tau < 1, \tag{9}$$

where τ measures the macro importance of the physical input.

Given the assumptions on the production of the physical and the intellectual inputs, the final output is given by a weighted CES:

$$Y = \left[\tau \left(\alpha_P K^{\rho_P} + (1-\alpha_P) L_P^{\rho_P}\right)^{\rho/\rho_P} + (1-\tau)\left(\alpha_I^{1-\rho_I} K_I^{\rho_I} + (1-\alpha_I)^{1-\rho_I} L_I^{\rho_I}\right)^{\theta_I \rho/\rho_I}\right]^{1/\rho}. \tag{10}$$

where $L_P = \alpha L$ and $L_I = (1-\alpha)L$.

This construction introduces modeling flexibility relative to the basic task setup of Acemoglu and Restrepo (2018) (also used in Jones and Liu (2024); Korinek and Suh (2024)), where the elasticity of substitution between tasks is closely linked to the substitution between labor and capital. Here, we have a macro elasticity of substitution between physical and intelligence tasks that can differ from the substitution patterns within the physical block and within the intelligence block. This allows us to meaningfully talk about how physical and intelligence production components behave differently and ultimately affect workers differently.

Scaling behavior Since the amount of AI capital K_I is increasing exponentially and very rapidly, it is important to understand how this scales up the intelligence input I, and ultimately output Y.

First, while the aggregate production function has constant returns to scale, we allow for diminishing returns to scale $\theta_I \leq 1$ in the production of the intelligence input I. This means that the scaling up of AI has more strongly diminishing returns at the aggregate level than in a case where the production function has constant returns to scale at all levels.

Second, the CES production function in P and I has the well-known property that output Y saturates when intelligence and physical outputs are complements ($\rho < 0$), and I goes to infinity while P is held fixed. This means that no matter how quickly AI capital K_I increases, the complementarity between physical and intelligence inputs means that output is bounded, and P becomes the bottleneck.

The complementarity between physical and intelligence inputs is a way to use a standard macro setting to represent insights from neuroscience and psychology that using more intelligence with fixed effectors cannot yield unbounded output growth.

3.3 Labor allocation and wages as automation progresses

We seek to answer two questions. As α_I, the share of automated tasks increases:

1. What is the share of labor allocated to physical vs. intellectual production?
2. What is the marginal productivity of labor, and hence the competitive wage?

In equilibrium, the marginal product of labor be equalized in the physical and intelligence inputs, so, under partial automation, we must have for optimality:

$$\frac{\partial F}{\partial P}\frac{dP}{dL_P} = \frac{\partial F}{\partial I}\frac{dI}{dL_I} \tag{11}$$

This determines the share of labor α^* that is allocated to the physical task. At the optimum, shifting an infinitesimal amount of labor from the intellectual block to the physical block yields zero net gain in output.

3.3.1 Macro blocks model: wages and labor allocation

We can derive explicitly the marginal productivity of labor in each block. By the chain rule, we find the competitive wage as the marginal productivity of labor:

$$\frac{\partial Y}{\partial L} = Y^{1-\rho}\Big[\tau\,\rho\,(1-\alpha_P)\,\alpha_P\,P^{1-\rho_P/\rho}\,L_P^{\rho_P-1} + (1-\tau)\,\theta_I\,\rho\,(1-\alpha_P)\,(1-\alpha_I)^{1-\rho_I}\,I^{1-\rho_I/(\theta_I\rho)}\,L_I^{\rho_I-1}\Big]. \tag{12}$$

As an aside, if wages are subject to a markdown, which is a realistic assumption (Azar and Marinescu, 2024), then the important question is whether the markdown differs between physical and intelligence tasks. If the wage markdown is markedly different, it would affect firms' decisions to allocate workers to physical vs. intelligence inputs. Markdowns would also obviously lower the actual wage received by workers.

We can then solve for α^*, which is the share of labor in physical tasks that equates the marginal productivity of labor in physical and intellectual tasks. Under the simplifying assumption that when α_I changes we treat the intermediate inputs P and I (and their own marginal products $\frac{dP}{dL}$, $\frac{dI}{dL}$) as approximately fixed, one obtains the closed-form:

$$\alpha^* = \Big[1 + \frac{1-\tau}{\tau} \frac{\theta_I \, (1-\alpha_I)^{1-\rho_I} \, L^{\rho_I - \rho_P} \, I^{1-\frac{\rho_I}{\theta_I \rho}}}{(1-\alpha_P) \, P^{1-\frac{\rho_P}{\rho}}}\Big]^{-\frac{1}{\rho_I}}. \tag{13}$$

The wage can be recovered for the optimal allocation by plugging α^* into the marginal product of labor in the physical task:

$$w = \tau \, Y^{1-\rho} \, \rho \, (1-\alpha_P) \, P^{1-\frac{\rho_P}{\rho}} \, (\alpha^* \, L)^{\rho_P - 1}. \tag{14}$$

As automation moves labor into the physical input production, the marginal productivity of labor in the physical block decreases, which has a negative impact on wages. But this negative effect on wages might be counteracted by an increase in the productivity of the intelligence input. We now discuss under what conditions wages may increase with automation.

3.3.2 Wage Dynamics as Automation Progresses in the Macro-Blocks model

We study how the competitive wage responds to an incremental rise in the automated share α_I, holding K and K_I fixed.

1. Full marginal-wage expression

By the chain rule on the two-block CES

$$Y = \Big[\tau P^\rho + (1-\tau) I^\rho\Big]^{1/\rho},$$

we have

$$\frac{dw}{d\alpha_I} = (A) + (B) + (C), \tag{15}$$

where

$$(A) = (1-\rho) \, Y^{-\rho} \, \underbrace{\frac{\partial Y}{\partial I}}_{F_I} \, \underbrace{\frac{dI}{d\alpha_I}}_{>0} \Big(\tfrac{dP}{dL} + \tfrac{dI}{dL}\Big) > 0, \tag{16}$$

$$(B) = Y^{1-\rho} \, \frac{d}{d\alpha_I}\Big(\tfrac{dP}{dL}\Big) < 0, \quad (C) = Y^{1-\rho} \, \frac{d}{d\alpha_I}\Big(\tfrac{dI}{dL}\Big) < 0. \tag{17}$$

(A) is the *scale-effect*: more automation raises I, which feeds into Y. (B) and (C) are *reallocation costs*: freed labor crowds respectively into P and residual I, depressing their marginal products.

Note that, once full automation is achieved, the intelligence block can be written as $I = K_I^{\theta_I}$, so the ρ_I parameter drops out. This means that the substitutability between intelligence tasks does not influence the level of output Y at full automation, nor the marginal productivity of labor at full automation.

2. Key approximations: "early-automation, abundant-AI"

We are interested in deriving the impact of parameters on the wage effect of automation $\frac{dw}{d\alpha_I}$. To do so, we make three simplifying assumptions:

(A1) *Abundant AI:* The stock of AI is so large that each additional automation transforms directly into more I with CES curvature. Concretely,

$$I(\alpha_I) \approx \left[\alpha_I^{1-\rho_I} K_I^{\rho_I}\right]^{\theta_I/\rho_I} = K_I^{\theta} \alpha_I^{\omega}, \quad \omega \equiv \frac{\theta_I (1-\rho_I)}{\rho_I},$$

hence

$$\frac{dI}{d\alpha_I} \approx \frac{\theta_I (1-\rho_I)}{\rho_I} \frac{I}{\alpha_I}.$$

(A2) *Negligible dilution:* In the "early-automation, abundant-AI" limit, α_I and total I rise almost proportionally, so their ratio α_I/I moves only second order in α_I. We therefore treat it as a constant normalized to 1: $\alpha_I/I \approx 1$, which eliminates a scale term that would otherwise carry through every derivative yet never affect signs.

(A3) *Small-shock linearization:* All second- and higher-order terms in $\Delta\boldsymbol{\alpha}_I$ are neglected, so $(\mathbf{B}) + (\mathbf{C})$ are evaluated to first order in ΔI.

Under condition A1, the scale effect (A) can be rewritten as:

$$(A) \approx (1-\rho)\, Y^{-\rho}\, F_I\, (F_P + F_I)\, \frac{\theta_I (1-\rho_I)}{\rho_I} \frac{I}{\alpha_I}. \tag{18}$$

(A) is positive because all of the terms are strictly positive. In particular, by assumption $\theta_I > 0$ and $0 < \rho_I < 1$.

The net effect of early automation on wages is positive whenever the positive scale effect (A) outweighs the negative reallocation drags (B) and (C) (equation 15). Under assumption **(A3)** (Small-shock linearization), we truncate $(B) + (C)$ at first order in $\Delta\alpha_I$. We also invoke **(A2)** (Negligible dilution) to set $I/\alpha_I \approx 1$. Under these assumptions the condition for early automation increasing wages simplifies to

$$\frac{\theta_I(1-\rho_I)}{\rho_I} > \frac{1}{(1-\rho)\dfrac{F_I}{F_P + F_I}}. \tag{19}$$

As we can see looking at the left-hand side of equation 19, wages initially rise when decreasing returns to scale are not too steep, i.e. θ_I is large enough. Looking at the right-hand side of the equation, we can see that wages are more likely to increase if the marginal productivity of the intelligence input F_I is large relative to the marginal productivity of the physical input F_P. The right-hand side also shows that wages are more likely to rise if the substitution between physical and intelligence inputs ρ is low (more complementarity).

To track the role of τ, the capital share in the physical input, note that:

$$\frac{F_I}{F_P + F_I} = \frac{(1-\tau)\, I^{\rho-1}}{\tau\, P^{\rho-1} + (1-\tau)\, I^{\rho-1}}.$$

Thus the share of capital in the aggregate production function, τ, enters indirectly by shifting the share of intelligence I in the marginal product.

Concretely, wages are more likely to increase in the low-automation abundant AI case if:

- $\theta_I \uparrow$: The intelligence block itself exhibits milder diminishing returns.
- $\rho_I \to 0$: Intelligence tasks become strong complements, so each additional bit of automation delivers large gains in I (see left-hand side of equation 19).
- $\rho \downarrow$: Physical and intelligence inputs are less substitutable in the top-level CES, so extra I has a larger marginal impact on output.
- $\tau \downarrow$: A smaller weight on the physical block raises the marginal-product share of I, making it easier for its scale-effect to dominate.

Existence of an interior turning point α_I^*. If the positive-wage condition in equation 19 holds at $\alpha_I = 0$, but eventually $(B) + (C)$ grow large enough to dominate (A), then by continuity there is a unique $\alpha_I^* \in (0, 1)$ solving

$$\frac{dw}{d\alpha_I} = 0.$$

Solving the equilibrium condition $F_I/(F_P + F_I) = \rho_I/[\theta_I(1-\rho_I)(1-\rho)]$ under our approximations yields

$$\alpha_I^* = \left[\underbrace{\frac{\rho_I\, \tau}{\theta_I\,(1-\rho_I)\,(1-\rho)\,(1-\tau)}}_{A} \times \underbrace{\left(\frac{\bar{P}}{K_I^{\theta_I}}\right)^{\rho-1}}_{B} \right]^{\frac{\rho_I}{(\theta_I(1-\rho_I))(\rho-1)}}. \tag{20}$$

where $\bar{P}$ is the physical-block marginal product held roughly constant.

Monotonic vs. hump-shaped wages. Depending on parameter values, the impact of automation on wages could be monotonic, or U-shaped, or hump-shaped.

Let's assume the effect is monotonically increasing or hump-shaped, and derive the impact of the parameters on the automation share where the wage peaks. In other terms, we want to ask if the parameters accelerate or delay a potential wage decline, assuming that the wage was initially increasing with automation.

Comparative-statics on the wage-peak α_I^*

Because $\rho - 1 < 0$, increasing any parameter that raises the bracket $A \times B$ in equation 20 will, via the negative exponent, decrease α_I^*. In particular (see appendix A for derivations):

- **Within–I scale θ_I:** A larger θ_I lowers A (denominator increases) and increases B (because B is raised to a negative power), acting in opposite directions. The first effect dominates under the abundant AI condition and shrinks the bracket—and since the exponent is negative, that raises α_I^*. As a result, higher θ_I delays the wage peak. Faster increasing returns in the intelligence block postpone the point at which the reallocation drag dominates.

- **Intelligence complementarity ρ_I:** Lowering ρ_I (stronger complementarity among intelligence tasks) both lowers A and increases the negative exponent on the bracket. Either way the net effect is to raise α_I^* so lower ρ_I delays the wage peak.

- **Cross-block elasticity ρ:** Reducing ρ (stronger complementarity across P and I) raises $(1-\rho)$ in the denominator of A, which lowers the bracket—and hence raises α_I^*. So lower ρ delays the wage peak.

- **Physical-share τ:** Raising τ raises the bracket directly (since $A \propto \tau$), and via the negative exponent decreases α_I^*. So higher τ accelerates the wage peak.

Economic intuition. Automation can increase wages whenever AI generates sufficiently large intelligence gains that feed effectively into output; otherwise wages decrease even at low levels of automation. Assuming wages increase at low levels of automation, then, if these intelligence gains remain strong throughout, wages rise all the way to full automation. If instead the reallocation drag eventually dominates, wages peak at $\alpha_I^* < 1$ and then decline. The parameters $\theta_I, \rho_I, \rho, \tau$ jointly determine which scenario obtains and where any peak occurs.

In Table 2, we summarize how parameters in the two-blocks model influence the impact of automation on wages: whether they make it more likely that wages rise initially if AI is abundant, and whether, assuming wages rise, the wage peak occurs later in the automation process. For relative physical scarcity, the table assumes that $\bar{P}/(K_I^{\theta}) < 1$, i.e. that there is little physical output relative to the amount of AI, consistent with the assumption of abundant AI. These comparative-statics are derived under the "early automation, abundant AI" approximation and require $0 < \rho_I < 1$. If $\rho_I < 0$, the above approximation no longer applies, but we can still solve the model numerically, which we will do in the next section.

Parameter increased	**Description**	**Wages rise initially?**	**Wage peak delayed?**
θ_I	Returns to scale in the intelligence block	Yes	Yes
ρ_I	Substitutability among intelligence tasks	No	No
ρ	Substitutability between physical and intelligence	No	No
τ	Weight on physical input in the aggregate CES	No	No

Table 2: **Parameter impacts on initial wage response and on the location of the hump-shaped peak α_I^*.**
See text for approximations needed to come to these predictions.

4 Simulations

In this section we will ask if the wage patterns predicted analytically under some simplifications survive when we calibrate the model with empirical parameter ranges.

The goal of these simulations is to predict the trajectory of wages as the automation of intelligence progresses, while holding everything else constant. Therefore, these simulations are most relevant to understand what happens in the short to medium run in the event of rapid automation, so that other factors do not adjust much. In the longer run, we can expect more adjustments, and in particular increases in physical capital. Nevertheless, the forces described here are still at play, and additional adjustments would have to undo the effects uncovered here.

4.1 Calibration

Using knowledge from the existing literature, we calibrate the macro blocks model (see Table 3). We calibrate share parameters α_P and τ using factor ratios and factor shares, as well as the substitution parameter. The factor share α is the payment to that factor divided by the payment to both factors, so for capital K and labor L with rental rate r and wage w, we have $\alpha_K = \frac{rK}{rK+wL}$. Given a CES in the two factors, we can use the fact that the factor payment is equal to its marginal productivity to derive, for the physical block,

$$\alpha_P = \frac{\alpha_K}{\alpha_K + (1-\alpha_K)\left(\frac{K_P}{L_P}\right)^{\rho_P}} \tag{21}$$

The analogous formula can be used to derive τ, the weight on the physical block in the aggregate CES.

First, we calibrate the physical block and hold constant the values of these parameters once we calibrate them. The capital share in the US economy is 42% in 2023 (U.S. Bureau of Labor Statistics, 2025a), and the capital/labor ratio in 2023 can be calculated as the ratio of the capital stock and the total labor compensation, which is about 5.8. Using eq. 21, we find that $\alpha_P = 0.7$. We set K_P so that K_P/L_P is equal to the overall capital/labor ratio in the economy of 5.8, given $L_P = 0.26$, reflecting the share of earnings in manual occupations in 2023, which is 0.26 (see Appendix B). The capital-labor substitution parameter ρ_P is set so that the elasticity of substitution $\sigma_P = 1/(1-\rho_P)$ is 0.6, corresponding to the middle of the range for US manufacturing (Oberfield and Raval, 2021).

Second, we calibrate parameters that we also allow to vary in different sweeps. The parameter determining the elasticity of substitution between AI capital and labor in the intelligence block ρ_I is set at baseline to reflect a high degree of substitutability, consistent with the idea that intelligence tasks are somewhat fungible with each other. Thus, for the baseline, we picked an estimate of the elasticity for somewhat similar goods: specifically, we used the smaller estimates for the elasticity of substitution among different physical goods in (Broda and Weinstein, 2006, Table IV, SITC-3, 1990-2001, median), which gives a baseline elasticity of substitution of 2.2. For the lower bound of ρ_I, we used the elasticity of substitution across occupations in Goos et al. (2014), which represents substitution between a broader range of tasks including both physical and intelligence tasks. For the baseline

returns to scale in intelligence, we use the 0.94 returns to scale in professional services from Figure 3 in (McAdam et al., 2024), as this is an industry that is dominated by intelligence tasks. There is no direct measurement of ρ, the substitution parameter between physical and intelligence inputs, so we set it to match the historical substitution between capital and labor in manufacturing, with the idea that capital is more like a physical input and labor more like an intelligence input. Practically, we set ρ so that the elasticity of substitution is at the middle of the range of the elasticity of substitution between capital and labor in US manufacturing, i.e. 0.6 (Oberfield and Raval, 2021). For the share of the physical block in output τ, we use equation 21, replacing with the variables relevant here, i.e. the share of earnings in manual occupations in 2023, which is 0.26 (see Appendix B), the ratio of manual to cognitive earnings in 2023 (to proxy for P/I), which is $0.26/(1-0.26) \approx 35\%$, as well as the baseline value for ρ. We obtain $\tau = 0.18$. Manual occupations are physical, and some non-manual occupations would also be physical under our definition of requiring in-person work, so the share of manual occupations in earnings is likely a understatement of the share of physical in output.

Having calibrated the model, we vary the share of automated tasks α_I from zero to one, simulating the process of progressively automating the intelligence tasks of the economy. For every setting, we find the share of labor allocated to the physical task α that equalizes the marginal productivity of labor in the physical and the intelligence blocks - the criterion needed for an equilibrium in the job market. We calculate this for each α_I up to $\alpha_I = 0.9$; the abundant AI condition 8 is no longer satisfied when α_I approaches 1, but it is still satisfied when $\alpha_I = 0.9$ for $K_I >= 9$ and $L_I = 1$. Then, at full automation for $\alpha_I = 1$, we set $\alpha = 1$ and calculate the output and wage based on these parameters instead of finding the optimal α.

Before moving on to the results of the simulations, it is worth noting that some intelligence tasks have already been automated as of 2025, so we should not think of α_I as starting at zero. As an approximation for the starting point, one may think about the percentage point decline in the share of routine cognitive tasks as a share of all cognitive tasks. Cognitive jobs do not perfectly overlap with intelligence jobs, because some cognitive jobs need to be done in person and are therefore "physical" in our classification. Nonetheless, this is a useful starting point. In 1976-1989, the share of routine jobs among all cognitive jobs was 49% (Cortes et al., 2020, Table 1), while this share was about 35% in 2018 (Cortes et al., 2020, Fig. 1), which represents a decline of 14 percentage points. This suggests that α might be around 14% in 2018, and somewhat higher today (2025).

4.2 Simulation results

Overall, the simulations illustrate that wages increase and then decrease with automation under plausible parameter values. When automation α_I increases, α – the share of labor allocated to the physical input – increases, and output Y increases, as expected given assumptions and eq. 8 and 13. Both total output and optimal allocation of workers to physical tasks increase with increasing automation (Fig. 1).

These dynamics raise the question of how wages are affected by increasing automation. Figure 2 plots the trajectory of wages as automation of intelligence tasks progresses, as a function of parameters. The orange line represents in each panel the trajectory of wages

Table 3: Simulation parameters and value ranges

Parameter	Description	Value(s)	Justification
Fixed baseline values (not calibrated)			
L	Total labor supply	1.0	Normalization
K_I	AI input level	9.0	Abundant AI
α_I	Automation share grid	0.01–0.99 (50 pts)	
Calibrated values held fixed			
K_P	Physical capital stock	1.5	For K_P/L_P to match 2023 aggregate capital/labor share of 5.8 when $L_P = 0.26$
α_P	Capital share in physical composite	0.7	To match 0.42 capital share in 2023 (U.S. Bureau of Labor Statistics, 2025a) given 2023 aggregate capital/labor share of 5.8 and baseline ρ_P.
ρ_P	Elasticity parameter in physical composite	-0.67	0.6 elasticity of substitution between capital and labor, middle of range in US manufacturing (Oberfield and Raval, 2021)
Calibrated values that vary			
ρ_I	Intelligence composite elasticity parameter	0.55 (baseline); -0.1, 0.55, 0.7 in sweep	Elasticity of substitution of 2.2 across imports in Broda and Weinstein (2006). Lower bound of -0.1 is 0.9 elasticity of substitution across occupations in Goos et al. (2014).
θ_I	Returns to scale in intelligence	0.94 (baseline); 0.1, 0.94, 1 in sweep	0.94: returns to scale in professional services (McAdam et al., 2024)
ρ	Aggregate output elasticity	-0.67 (baseline); -1, -0.67, -0.1 in sweep	0.6 elasticity of substitution between capital and labor, middle of range in US manufacturing (Oberfield and Raval, 2021)
τ	Share physical in output	0.18 (baseline); 0.15, 0.24, 0.4 in sweep	To match the share of earnings in manual occupations, given ratio of manual to cognitive earnings, and baseline ρ

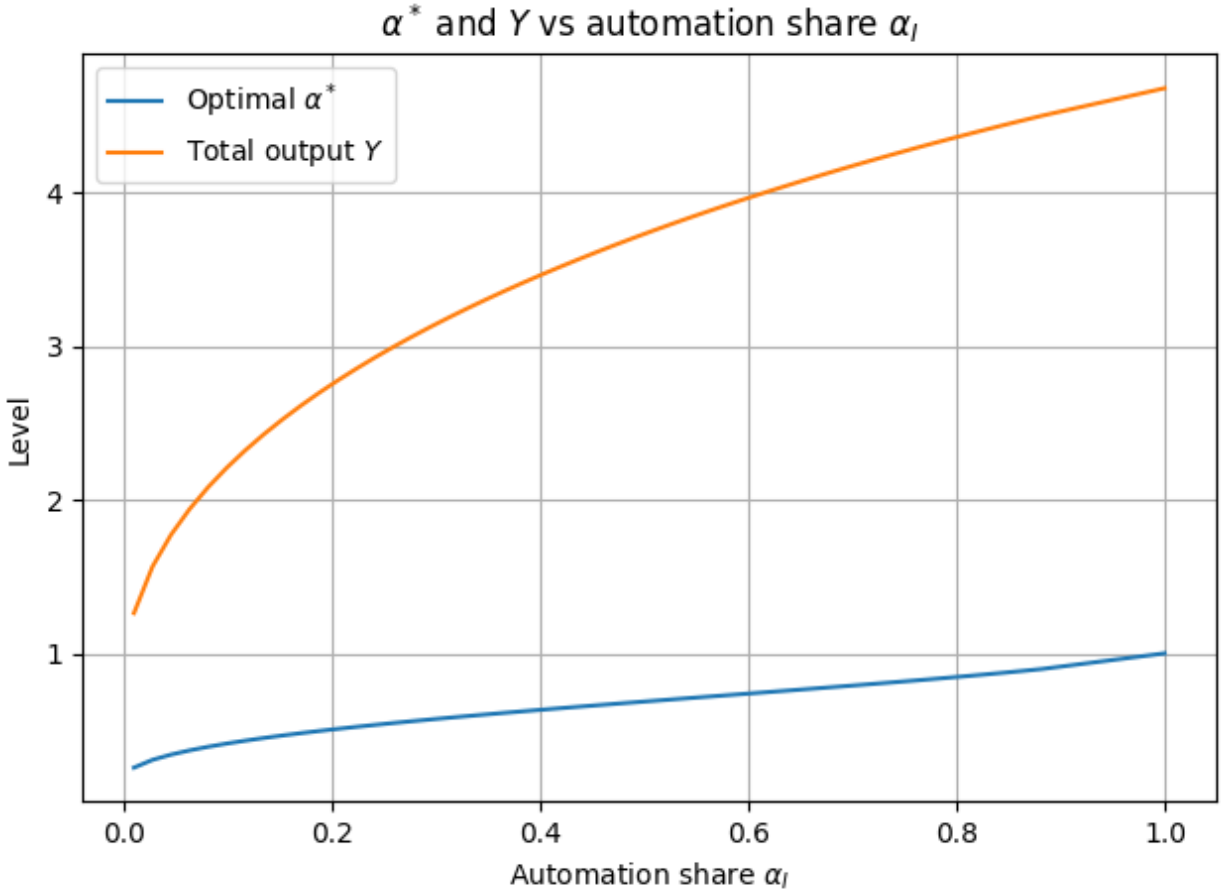


Figure 1: Output and the share of labor allocated to physical tasks increase with automation.

for the baseline set of parameters (note that the scale of the plot varies across panels). In the calibrated baseline case, wages increase and then decrease with automation, a trajectory that is frequently anticipated by many AI scientists. For example, this CIO report predicts short-term wage growth, while other works, such as comments by OpenAI CEO Sam Altman, suggest that the era of human work may be coming to an end. Across parameter settings, we can find situations where wages largely decline or increase. The exact trajectories of wages during the transition to intelligence automation will be heavily affected by the underlying substitution parameters.

The substitutability between different intelligence tasks, ρ_I should be one key ingredient. Without easy substitution, human intelligence activity in the non-automatable tasks keeps them employed and wages high. Variations in this parameter have a dramatic impact (Fig. 2 A). As ρ_I decreases, so that intelligence tasks become less substitutable, we see a stark hump-shaped pattern emerge: wages increase with low levels of automation, and then decrease, and the inflection point happens later than in the baseline, as predicted in Table 2 (the lowest ρ_I we examine in our sweep is negative, but the wage trajectory looks very similar for $\rho_I = 0.1$). The wage at full automation is not sensitive to ρ_I, because, as noted above, ρ_I drops out from the production function once automation is complete.

We have plausibly already automated some of the intelligence tasks, and average wages increased. Therefore, there is a risk of a strong decline in the future if intelligence tasks are relatively complementary. Essentially, more complementarity in intelligence tasks (lower ρ_I) makes the rise steeper but the downfall in wages worse as full automation approaches.

We may also believe that decreasing returns to scale in the intelligence production function may have an effect. Indeed, returns to scale do play a role, which interestingly only minimally interacts with the automation share (Fig. 2B). As decreasing returns to scale become more pronounced (lower θ_I), wages are more likely to decline initially (consistent with theoretical predictions in Table 2), but the shape of the path is also flatter. More decreasing returns to scale in intelligence almost erase the impact of automation on wages.

We may also expect the top-level substitution parameter ρ to be a main driver of wage effects. Similarly to the ρ_I substitution parameter, this parameter markedly changes the trajectory of wages as automation increases (Fig. 2C). As physical and intelligence inputs become less substitutable (smaller ρ), wages are less likely to decline with the automation process, again as was shown above in Table 2. There are dramatic differences in the wage level at full automation: wages are higher than without automation for $\rho = -1$, but they are a lot lower with $\rho = -0.1$. Therefore, a high level of substitution between physical and intelligence inputs drives a large loss for wages at full automation.

Finally, we examine the impact of the physical input share in the upper level CES between physical and intelligence inputs. This parameter strongly influences the impact of automation on wages (Fig. 2D). A larger weight on the physical input τ delays any decline in wages, which is not consistent with the approximation-based theoretical predictions in Table 2. For $\tau = 0.35$, wages monotonically increase with automation, while for $\tau = 0.1$ or $\tau = 0.18$, we see a large decrease in wages as we approach full automation. The wage outcome is extremely sensitive to the share of the physical block, and this share may decrease as automation progresses, which would yield negative wage effects.

Overall, our simulations mostly confirm the theoretical patterns predicted in Table 2; this was not a completely foregone conclusion since we relied on some approximations to

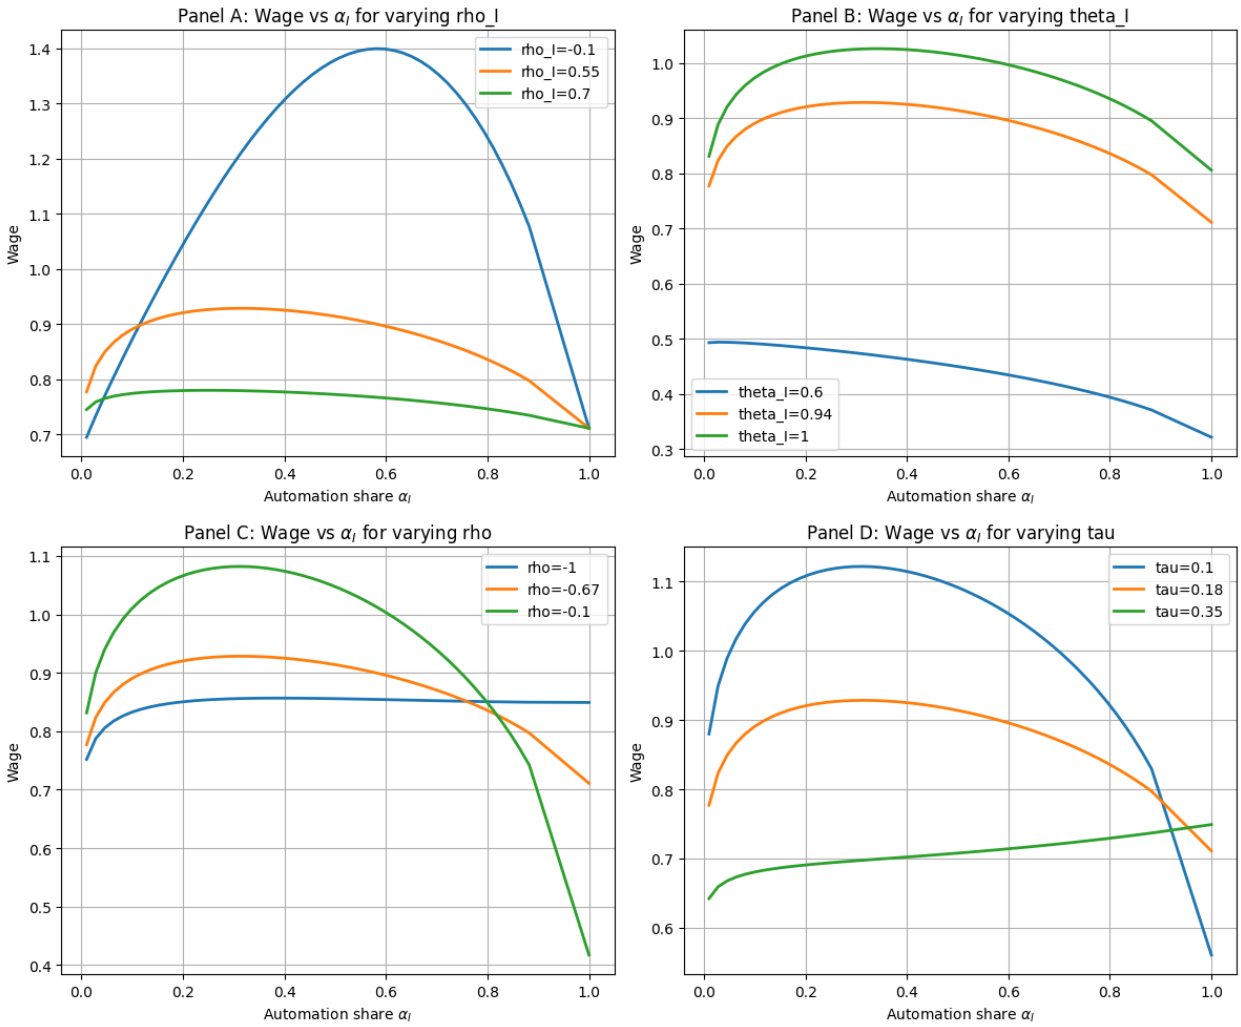


Figure 2: Wage effect of automation: impact of parameters

obtain our predictions. We further learn about the quantitative effects of these parameters on wages as automation progresses, and as we calibrate parameters based on the empirical evidence. One interesting pattern is that the impact of automation on wages tends to be very small if we have strongly decreasing returns in intelligence (low θ_I). Wages increase and then decrease with automation in the baseline case. The decrease in wages at high levels of automation is greater when there is low substitutability between intelligence tasks (low ρ_I), high substitutability between the physical and the intelligence inputs (high ρ), and a low share for the physical component (low τ).

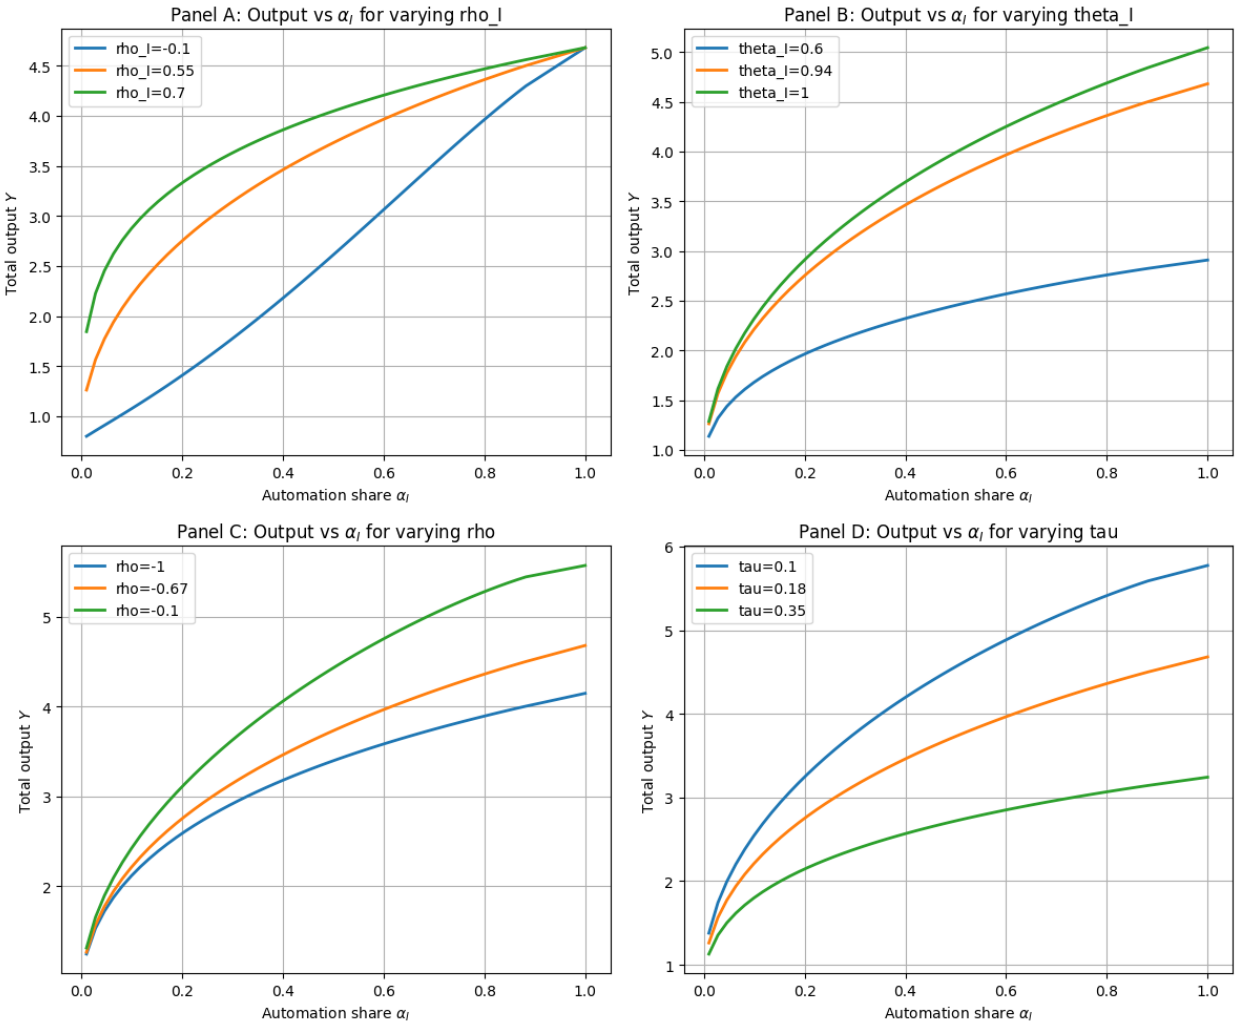


Figure 3: Output effect of automation: impact of parameters

Total output will matter on top of the wages. As such, we analyze in our model how automation affects total output (Fig. 3). As we should expect given the abundant AI assumption, automation always increases output, and considerably so. Higher substitutatbility between intelligence and physical inputs magnifies the positive effect of automation on output (Panel C) but has a negative effect on wages (compare panels C in Figure 3 vs. Figure 2). On the other hand, a higher share for the physical input reduces the output impact of

automation while it leads to a more positive impact of automation on wages (compare panels D in Figure 3 vs. Figure 2). In other terms, there is a trade-off between wages and output as we increase ρ the substitutability between the physical and intelligence inputs, or when we increase the share τ of the physical input. For wages, it's best to *protect* the physical sector where workers end up working more and more often as automation progresses: this means more complementarity between the physical and the intelligence inputs, and a higher share for the physical sector. This suggests a potential social conflict over the role of the physical sector in the economy when wage and output increases benefit different actors.

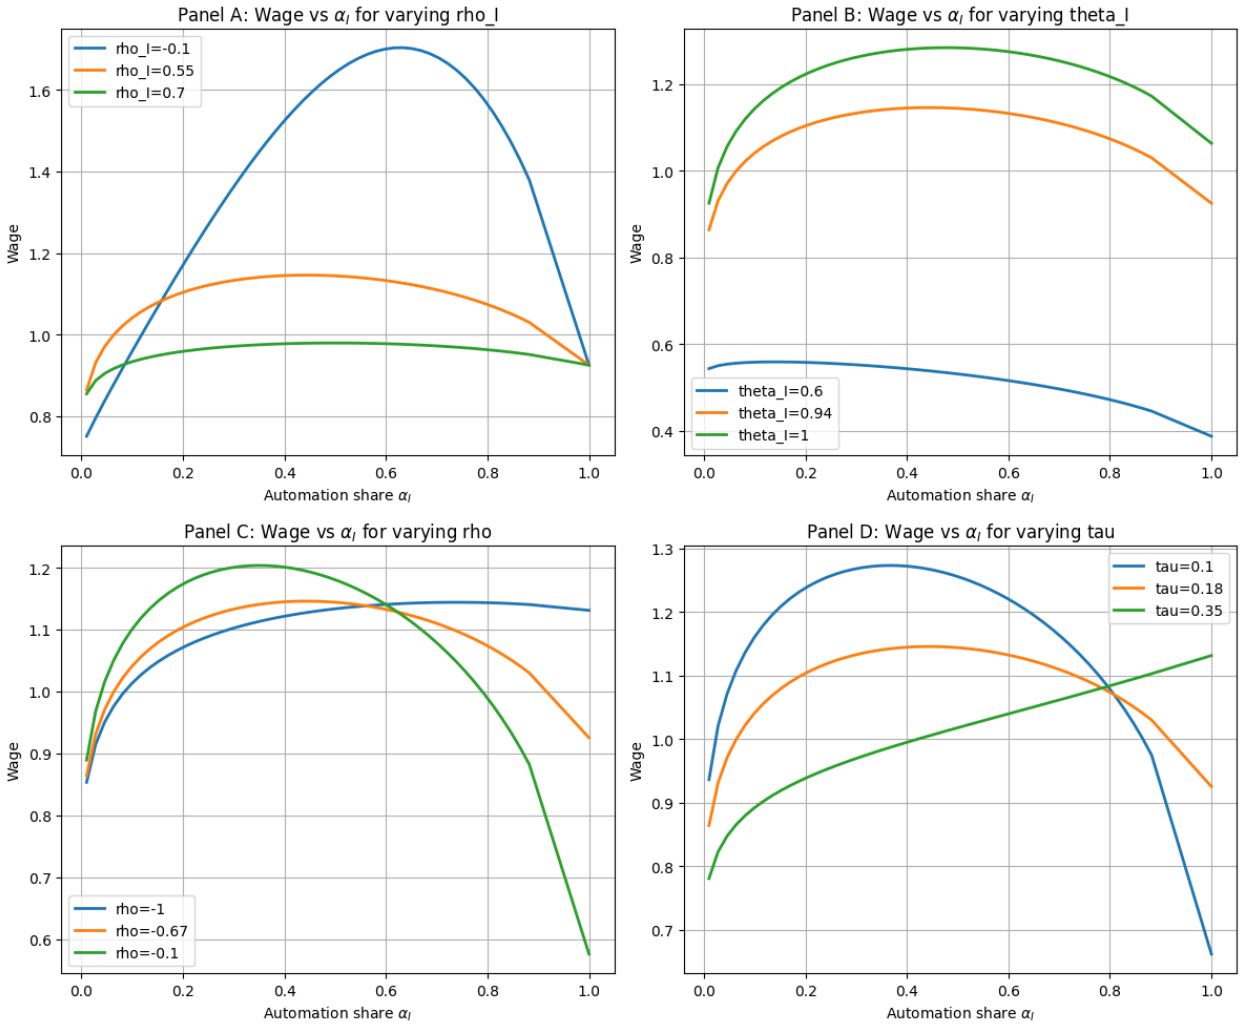


Figure 4: Wage effect of automation with high physical capital: impact of parameters

Given the crucial role of the physical sector, increasing capital in the physical sector can benefit wages because it reduces the negative productivity effects of labor crowding into the physical sector as automation progresses. This is illustrated in Figure 4, which uses the parameters from Table 3, except that physical capital K is set to 3 instead of 1.5. With a higher level of physical capital, wages increase more with automation: the maximum wage is higher and the initial increase is steeper than in Figure 2. If physical

capital were endogenously determined, investments would likely increase because returns to capital increase when there are more workers in the physical sector. Overall, increasing physical capital share can act as a stabilizer for wages.

We now turn to the longer-run effect of AI capital on wages, after all intelligence tasks have been automated. For this simulation, we assume a fixed 5 years for the full automation of intelligence tasks, during which α_I goes from 0 to 1. We further assume that AI capital K_I grows exponentially. All other parameters are as in the baseline, unless otherwise specified. In a first scenario of fast AI growth, AI doubles every 6 months. In this case, wages rise quickly and then saturate (Figure 5). The saturation is driven by the complementarity between physical and intelligence tasks, as our baseline has $\rho < 0$: you cannot add infinitely more AI and get unbounded output growth when the physical input is held fixed by fixed labor and fixed physical capital. In a second scenario with slower AI growth, AI doubles every 24 months. In this case, wages increase more slowly, but do not decline, a difference from the baseline analysis with fixed AI: the absence of a wage decline is explained by the fact that AI increases over time, which has positive effects on wages that offset the negative wage effects observed under fixed AI. Finally, in a third scenario, we continue assuming slower AI growth, but also make physical and intelligence inputs much more substitutable, with $\rho = 0.25$ (corresponding to a 1.33 elasticity of substitution). In this case, we observe a wage decrease before full automation of intelligence tasks, after which wages slowly grow again. Importantly, in this scenario with $\rho > 0$, there is no intelligence saturation, so in the long run adding more AI always keeps increasing wages. To sum up, as AI increases exponentially, wages saturate when physical and intelligence inputs are gross complements, but can grow without bounds when they are substitutes. However, greater substitutability of intelligence and physical inputs makes it more likely that wages decline as we approach full automation of intelligence tasks. Thus, there is a short-run vs. long-run trade-off in terms of the wage effects of substitutability between physical and intelligence: in the short run, greater substitutability can hurt wages as automation progresses, but in the long run it allows for unbounded wage growth, assuming AI also expands without bounds.

5 Discussion

Our family of models articulates how intelligence saturation may fundamentally limit the extent to which AI can transform the economy. The key insight is that intelligence and physical inputs are not perfect substitutes at the macro level, and AI's primary role is in enhancing intelligence production rather than physical production. Replacing some labor with AI is easy. Replacing more gets harder the more is already replaced, the old elasticity story. This insight has several important implications for understanding the impact of AI on the future of the economy.

5.1 Intelligence Saturation and the Non-Singularity

The concept of intelligence saturation fundamentally challenges the singularity narrative popularized by some AI researchers. The singularity view implicitly assumes unbounded returns to intelligence - that increasing intelligence inputs will lead to exponentially grow-

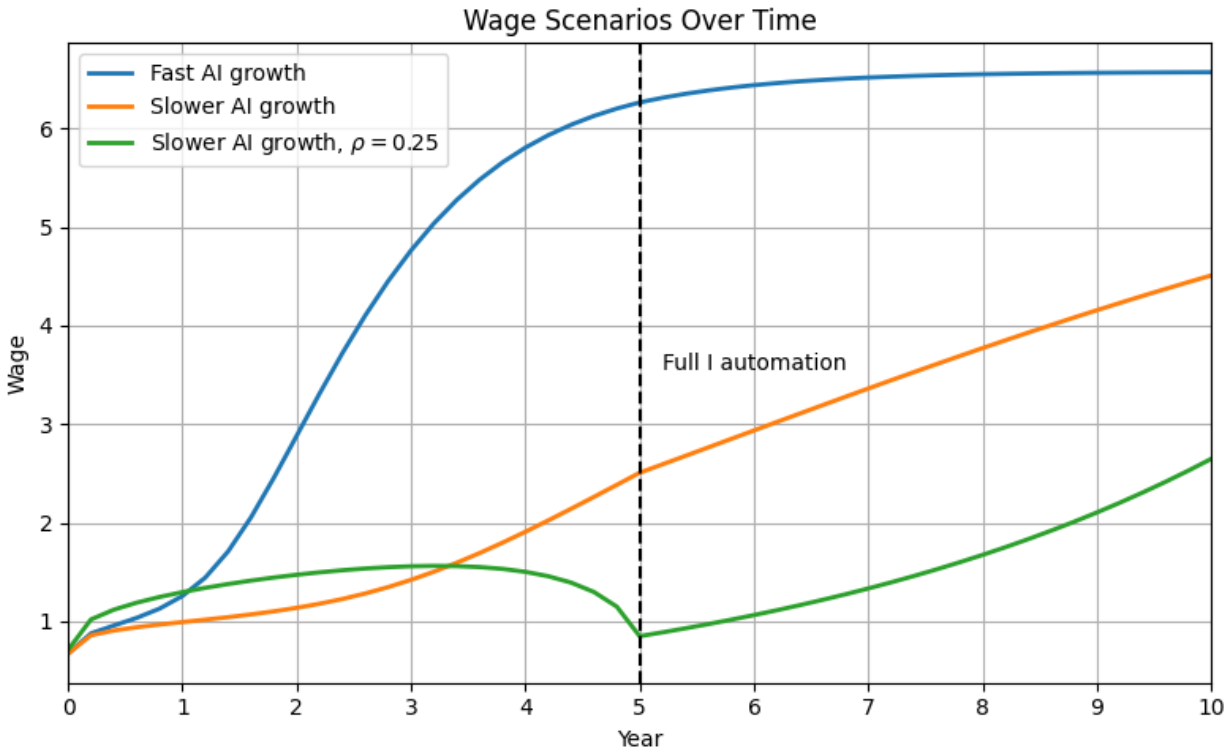


Figure 5: Wage effect of automation with exponentially increasing AI capital

ing productivity. Our model shows why this assumption fails: the physical world imposes constraints that cannot be overcome by intelligence alone.

For any given set of physical inputs, there exists a saturation point beyond which additional intelligence yields negligible returns. You cannot build a car with pure intelligence; you need steel, rubber, energy, and someone to fix the mess when something goes wrong. Simulations show that strong decreasing returns to intelligence lead to limited effects of automation on wages and output, even as labor optimally reallocates to the physical sector in response to automation. When physical and intelligence inputs are complementary, no amount of intelligence can remove the need for physical inputs or make them infinitely productive. The extremely rapid exponential growth of intelligence technology thus does not fundamentally alter the dynamics of growth.

5.2 Labor Market Implications Through Transition and Beyond

A key insight from our analysis is that labor allocation responds endogenously to automation. As α_I and K_I increase (more intelligence tasks become automated with more intelligence capital), labor shifts toward physical production according to optimal allocation. This reallocation itself depresses wages by increasing the labor supply in physical production relative to demand, but at the same time abundant AI in the intelligence sector makes physical workers more effective and thus boosts wages, so total wage effects depend on the balance of these two factors.

Our model predicts a hump-shaped pattern for wages for most calibrations, including our baseline calibration. Under some parameters, wages can decrease more with automation, in particular if there is high substitutability between the physical and the intelligence input. Once full automation of intelligence tasks is achieved, wages definitely grow again as the economy benefits from continuing investments in AI.

This wage trajectory differs from the extreme techno-optimist view where wages go to zero while productivity potentially goes to infinity. It also differs from the extreme pessimist view (which predicts permanently depressed wages). Instead, our model suggests a non-monotonic trajectory, with a sharp decrease in wages in later stages of automation, in particular when intelligence tasks are more complementary. Therefore, it is important to remain cautious: even as automation first increase wages, it can eventually lead to strong wage declines (also see Korinek and Suh (2024)). The substitution between intelligence tasks is critical and therefore should be studied more empirically, especially as this parameter may change with automation. Indeed, as automation progresses, it seems likely that the substitutability between still non-automated tasks and automated tasks decreases, in particular if there is a tail of intelligence tasks where labor has particularly strong advantages.

Given that automation can have a non-monotonic impact on wages, it is important to know where we are in the path to automation and what wage effects are already visible. Hampole et al. (2025) show that labor demand indeed decreases for jobs exposed to AI, but overall labor demand increases, leading to muted overall effects. These offsetting effects are consistent with the mechanisms in our model. Their estimates imply that the employment share of overall less exposed occupations like the physical "Food preparation and serving" declines because firms that do not use AI grow less. This suggests, through the lens of our model, a possible decline in τ, the share of physical in the economy, which could have adverse effects on wages. On the other hand, increasing AI share may well increase τ as increasing intelligence saturation lowers the gains to further AI investment. As such, the physical share τ both matters deeply for predictions and its temporal development is hard to predict.

The role of the physical sector is critical because as automation progresses, an increasing share of labor optimally reallocates to the physical sector. There is a trade-off between wages and output as a higher share of the economy in the physical sector and a low substituability between the physical and the intelligence inputs leads to more favorable automation effects for wages but *less* favorable automation effects for output. This is further affected by the fact that scaling in physical tasks is much slower than scaling in intellectual tasks. The most interesting parameter is perhaps ρ the substitution parameter between the physical and the intelligence inputs: if the goal is to sustain high wages while automation progresses, this parameter should stay low, but if the goal is to increase output the parameter should be increased. Concretely, this means that maximizing output may mean promoting virtual substitutes to physical goods and in person activities, which would then contribute to decreasing wages.

The substitution parameter ρ also fundamentally shapes intelligence saturation when other factors are held fixed: when $\rho < 0$, there is intelligence saturation, and the lower ρ, the lower the maximum wage and the quicker this maximum wage is achieved. When $\rho > 0$, i.e. when physical and intelligence inputs are gross substitutes, there is no intelligence saturation and wages can continue to grow forever as AI increases, but the wage trajectory during automation is typically less favorable.

5.3 AI, innovation and wages

AI leads to innovation that can itself spur job and wage growth in line with past introductions of general purpose technologies. This process of innovation is captured implicitly in two ways in our macro model.

First, our model can capture the creation of new labor-intensive intelligence tasks, a reinstatement effect (Acemoglu and Restrepo, 2019). Specifically, the creation of new tasks would manifest through a lower automation parameter α_I in the intelligence block, because the share of labor in intelligence tasks $1 - \alpha_I$ increases when new labor-intensive tasks are added (Acemoglu and Restrepo, 2019). Thus, negative wage effects of increasing automation α_I could be countered by the creation of new labor-intensive intelligence tasks.

Second, the progress of AI can be represented through innovation embodied in capital, by adding more artificial intelligence capital K_I. This increases the wages of intelligence workers for the same reason that adding more capital in a CES production function increases wages, i.e. because labor benefits from having more capital to work with. Adding more AI also increases the wages of physical workers because it increases the intelligence input, which is beneficial for the physical input to which physical workers contribute.

It is the process of automation by itself, i.e. the replacement of workers with AI, that can lead to negative wage effects if other positive effects are not strong enough to counteract it. While we do not explicitly model innovation, innovation effects are implicitly captured in the evolution of α_I and K_I, and can be interpreted as the aggregate outcome of innovation at the micro level.

5.4 Factors Moderating the Impact of AI

Several additional factors may moderate AI's economic impact beyond what our model captures. First, there may be intelligence tasks that remain persistently difficult to automate - tasks that require forms of intelligence where AI's capabilities plateau. Our modeling can accommodate this by capping α_I below 1, reflecting a frontier of non-automatable intelligence work.

Second, the rate of AI adoption depends on more than just technical capability - it requires investments, training, complementary innovations, and overcoming institutional inertia. These friction factors explain why technological transitions typically occur more gradually than pure technical capabilities would suggest.

Finally, a lack of competition among AI providers may limit AI development and deployment, or lead to strategic choices that diverge from socially optimal development paths.

5.5 Model Limitations and Future Research

Our models make several simplifying assumptions that should be acknowledged. First, we assume intelligence saturation is universal across all production activities. If there exist many domains where returns to intelligence do not diminish (or diminish very slowly), our conclusions about long-term growth would need modification.

Second, our models do not explicitly account for the role of human preferences and values in shaping the economic impact of AI. Even with intelligence saturation, societal choices

about AI governance, distributive justice, and work organization will significantly influence realized outcomes.

Third, our models do not explicitly incorporate reallocation costs for labor, while in practice these costs tend to be substantial. As a result, even if automation ultimately increases average wages, there can be significant wage declines in the short to medium run as workers must find new jobs and adapt their skills to these new jobs.

Fourth, because our model assumes labor is homogeneous, it cannot discuss disparate impacts of AI for different groups of workers. Just like prior waves of technological change (Autor and Dorn, 2013; Kogan et al., 2023), the newest AI technology may lead to wage increases for some workers and wage decreases for others depending on their sectors of activity and whether technology complements or substitutes for their work. It is worth noting that disparate impacts of automation across workers can be magnified in the presence of reallocation costs, which, as mentioned before, we do not explicitly model.

Future research could address these limitations by incorporating heterogeneous task structures with varying substitution elasticities, exploring dynamic adjustment processes in greater detail, and incorporating endogenous technological change in physical production capabilities. Additionally, enriching the models with evidence from recent AI deployments could improve their empirical grounding.

5.6 Policy Implications

Our analysis suggests several policy implications. First, policies that slow automation while increasing capital investment in physical sectors could help smooth the transition by reducing negative wage effects. This does not mean halting AI development, but rather pacing its deployment to allow for complementary investments in physical capital. A tax on the deployment of AI that replaces workers in intellectual tasks might be helpful to maintain high wages, but it is not clear how one could effectively target this technology. As our model outlines, replacing workers tends to reduce wages, but increasing the amount of AI used tends to increase wages; this makes the targeting of any AI tax extra tricky, even if AI could be clearly distinguished from other technologies. A subsidy for investment in capital used in physical inputs is likely easier to target: this could include for example physical investments in construction, and in the hospitality industry.

Second, redistributive policies may be useful during the transition period when wages are depressed by automation. The specifics of these policies depend on political and design considerations beyond our model's scope, but the model clearly indicates a period of adjustment where many workers face significant economic pressure. One possible model is wage insurance Hyman et al. (2024), which would provide income support to workers displaced by AI that find jobs at lower wages.

6 Conclusion

We have presented a framework for analyzing the economic impact of AI that reconciles insights from both economics and computer science. Our models demonstrate how intelligence

saturation - the principle that returns to intelligence diminish when physical inputs are held constant - fundamentally constrains AI's long-term economic impact.

The key insight of our analysis is the distinction between intelligence and physical inputs, coupled with the observation that labor is more substitutable with AI in intelligence tasks than with capital in physical tasks. These two features drive the dynamics of our models and lead to our main conclusions: (1) wages may have a non-monotonic path as automation progresses (2) after full automation of intelligence tasks, wages can grow again but remain limited by intelligence saturation, and (3) the ultimate constraint on economic growth shifts from intelligence to physical production.

We calibrate a CES macro-blocks model built on a CES physical block (capital–labor), and a CES intelligence block that aggregates many tasks. Under our baseline calibration, rising automation of intelligence tasks increases and then decreases wages, as the labor reallocation drags eventually outweigh the positive impacts on the production of intelligence. This represents a cautionary tale that early days wage gains should not be taken for granted. Varying parameters away from the baseline, we learn three main lessons. First, starker decreasing returns to intelligence dampen the wage effects of automation. Second, less substitutability between intelligence tasks can lead to wages strongly increasing and then steeply decreasing with automation. Third, because automation pushes labor toward the physical block, the elasticity of substitution between blocks is pivotal: greater substitutability depresses wages further yet boosts output, while a larger physical share cushions wages but limits aggregate gains. Further, this elasticity governs whether wages and output saturate as AI increases: when physical and intelligence inputs are gross complements ($\rho < 0$), there is intelligence saturation. The effects of automation depend on key parameters in a subtle way, making predictions parameter-dependent.

Overall, these findings contradict the singularity narrative that posits unbounded returns to intelligence. No matter how powerful AI becomes, it cannot escape the fundamental constraints imposed by the physical world. A car cannot be built by intelligence alone; a building cannot be constructed with pure computation; a meal cannot be cooked with algorithms alone.

Our analysis suggests that AI represents a powerful general-purpose technology that will significantly reshape the economy, but not one that fundamentally alters the laws of economics or physics. The economy after AI will still be constrained by scarcity, still require physical production, and still operate according to principles of marginal returns and substitution that economists have long understood.

This perspective provides a middle ground between the extremes that have dominated much public discourse about AI. Neither the techno-utopian vision of post-scarcity abundance nor the dystopian fear of permanent labor displacement accurately captures what our model predicts. Instead, we foresee a significant but ultimately bounded transformation, with a potentially challenging transition period followed by a new equilibrium where physical production becomes the limiting factor.

Policy can focus on managing this transition while ensuring that the benefits of AI are broadly shared. By recognizing the fundamental constraints on AI's economic impact, we can develop more realistic expectations and more effective strategies for harnessing AI's potential while mitigating its disruptive effects.

The future path of the economy depends not only on technological possibilities but also

on societal choices. We retain substantial agency in shaping how these technologies are developed and deployed, and AI itself can help us develop scenarios and possible policy responses.

In summary, our paper contributes to the understanding of AI's economic impact by providing a framework that combines insights from economics and computer science, emphasizes the fundamental distinction between physical and intelligence inputs, and demonstrates how intelligence saturation constrains AI's transformative potential. This framework helps resolve apparent contradictions between optimistic and pessimistic AI narratives and provides guidance for policy responses that can smooth the transition to an AI-abundant economy.

References

Acemoglu D, Autor D. 2011. Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings. In Card OAaD (ed.) *Handbook of Labor Economics*, volume Volume 4, Part B. Elsevier. ISBN 1573-4463, 1043–1171.
URL http://www.sciencedirect.com/science/article/pii/S0169721811024105

Acemoglu D, Kong F, Restrepo P. 2024. Tasks At Work: Comparative Advantage, Technology and Labor Demand.
URL https://www.nber.org/papers/w32872

Acemoglu D, Restrepo P. 2018. Low-Skill and High-Skill Automation. *Journal of Human Capital* **12**: 204–232. ISSN 1932-8575, 1932-8664.
URL https://www.journals.uchicago.edu/doi/10.1086/697242

Acemoglu D, Restrepo P. 2019. Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives* **33**: 3–30. ISSN 0895-3309.
URL https://www.aeaweb.org/articles?id=10.1257%2Fjep.33.2.3

ARK Investment Management LLC. 2021. Industrial robot costs could fall 65% by 2025, boosting automation. Accessed: 2025-05-07.
URL https://www.ark-invest.com/articles/analyst-research/industrial-robot-costs

Autor DH, Dorn D. 2013. The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review* **103**: 1553–1597. ISSN 0002-8282.
URL http://pubs.aeaweb.org/doi/10.1257/aer.103.5.1553

Autor DH, Levy F, Murnane RJ. 2003. The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics* **118**: 1279–1333. Accessed July 8, 2025.
URL https://www.jstor.org/stable/25053940

Azar J, Marinescu I. 2024. Chapter 10 - Monopsony power in the labor market. In Dustmann C, Lemieux T (eds.) *Handbook of Labor Economics*, volume 5. Elsevier, 761–827.
URL https://www.sciencedirect.com/science/article/pii/S1573446324000099

Barrero JM, Bloom N, Davis SJ. 2021. Why Working from Home Will Stick.
URL https://www.nber.org/papers/w28731

Bloom N, Jones CI, Van Reenen J, Webb M. 2020. Are Ideas Getting Harder to Find? *American Economic Review* **110**: 1104–1144. ISSN 0002-8282.
URL https://www.aeaweb.org/articles?id=10.1257/aer.20180338

Broda C, Weinstein DE. 2006. Globalization and the Gains From Variety*. *The Quarterly Journal of Economics* **121**: 541–585. ISSN 0033-5533.
URL https://doi.org/10.1162/qjec.2006.121.2.541

Brown TB, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, et al. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems* **33**: 1877–1901.
URL https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf

Cortes GM, Jaimovich N, Nekarda CJ, Siu HE. 2020. The dynamics of disappearing routine jobs: A flows approach. *Labour Economics* **65**: 101823. ISSN 0927-5371.
URL https://www.sciencedirect.com/science/article/pii/S0927537120300270

Cortés-Albornoz MC, Ramírez-Guerrero S, García-Guáqueta DP, Vélez-Van-Meerbeke A, Talero-Gutiérrez C. 2023. Effects of remote learning during COVID-19 lockdown on children's learning abilities and school performance: A systematic review. *International Journal of Educational Development* **101**: 102835. ISSN 0738-0593.
URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10266495/

DeepLearningAI. 2024. Falling llm token prices and what they mean for ai companies. Accessed: 2025-05-07.
URL https://www.deeplearning.ai/the-batch/falling-llm-token-prices-and-what-they-mean

Goos M, Manning A, Salomons A. 2014. Explaining Job Polarization: Routine-Biased Technological Change and Offshoring. *American Economic Review* **104**: 2509–2526. ISSN 0002-8282.
URL https://www.aeaweb.org/articles?id=10.1257/aer.104.8.2509

Guo D, Yang D, Zhang H, Song J, Zhang R, Xu R, Zhu Q, Ma S, Wang P, Bi X, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948* .

Hampole M, Papanikolaou D, Schmidt LD, Seegmiller B. 2025. Artificial Intelligence and the Labor Market.
URL https://www.nber.org/papers/w33509

Hernandez D, Brown TB. 2020. Measuring the algorithmic efficiency of neural networks. *arXiv preprint arXiv:2005.04305* .
URL https://arxiv.org/abs/2005.04305

Hyman BG, Kovak BK, Leive A. 2024. Wage Insurance for Displaced Workers.
URL https://www.nber.org/papers/w32464

Jack R, Halloran C, Okun J, Oster E. 2023. Pandemic Schooling Mode and Student Test Scores: Evidence from US School Districts. *American Economic Review: Insights* **5**: 173–190. ISSN 2640-205X.
URL https://www.aeaweb.org/articles?id=10.1257/aeri.20210748

Jones BF. 2009. The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder? *The Review of Economic Studies* **76**: 283–317. ISSN 0034-6527.
URL https://doi.org/10.1111/j.1467-937X.2008.00531.x

Jones BF, Liu X. 2024. A Framework for Economic Growth with Capital-Embodied Technical Change. *American Economic Review* **114**: 1448–1487. ISSN 0002-8282.
URL https://www.aeaweb.org/articles?id=10.1257/aer.20221180

Kelly B, Papanikolaou D, Seru A, Taddy M. 2021. Measuring Technological Innovation over the Long Run. *American Economic Review: Insights* **3**: 303–320. ISSN 2640-205X, 2640-2068.
URL https://pubs.aeaweb.org/doi/10.1257/aeri.20190499

Kogan L, Papanikolaou D, Schmidt L, Seegmiller B. 2023. Technology and Labor Displacement: Evidence from Linking Patents with Worker-Level Data.
URL https://papers.ssrn.com/abstract=3585676

Koomey JG, Berard S, Sanchez M, Wong H. 2011. Implications of historical trends in the electrical efficiency of computing. *IEEE Annals of the History of Computing* **33**: 46–54.

Korinek A, Suh D. 2024. Scenarios for the Transition to AGI.
URL https://www.nber.org/papers/w32255

Kurzweil R. 2005. *The Singularity Is Near: When Humans Transcend Biology*. New York: Viking.

McAdam P, Meinen P, Papageorgiou C, Schulte P. 2024. Returns to Scale: New Evidence from Administrative Firm-Level Data.
URL https://www.ssrn.com/abstract=4886095

Nordhaus WD. 2007. Two centuries of productivity growth in computing. *The Journal of Economic History* **67**: 128–159.

Oberfield E, Raval D. 2021. Micro Data and Macro Technology. *Econometrica* **89**: 703–732. ISSN 0012-9682. Publisher: Wiley-Blackwell Publishing Ltd.
URL https://collaborate.princeton.edu/en/publications/micro-data-and-macro-technology

Park M, Leahey E, Funk RJ. 2023. Papers and patents are becoming less disruptive over time. *Nature* **613**: 138–144. ISSN 1476-4687. Publisher: Nature Publishing Group. URL https://www.nature.com/articles/s41586-022-05543-x

Ray SC, Chambers RG, Kumbhakar SC (eds.) . 2022. *Handbook of Production Economics*. Singapore: Springer Nature Singapore. ISBN 978-981-10-3454-1 978-981-10-3455-8. URL https://link.springer.com/10.1007/978-981-10-3455-8

US Bureau of Labor Statistics. 2001. *Using the Occupational Classification System Manual (OCSM) : U.S. Bureau of Labor Statistics*. Accessed June 30, 2025. URL https://www.bls.gov/ocs/additional-resources/census-job-titles-see-ocsm.htm

US Bureau of Labor Statistics. 2023. *May 2023 National Occupational Employment and Wage Estimates*. Accessed June 30, 2025. URL https://www.bls.gov/oes/2023/may/oes_nat.htm#00-0000

US Bureau of Labor Statistics. 2025a. Private Business Sector: Capital Share [MPU4900131]. Retrieved from FRED, Federal Reserve Bank of St. Louis. URL https://fred.stlouisfed.org/series/MPU4900131

US Bureau of Labor Statistics. 2025b. Producer price index by industry: Construction machinery manufacturing: Power cranes, draglines, and shovels (excavators) (including surface mining equipment) (excluding parts). Index Dec 1999=100, Monthly, Not Seasonally Adjusted. URL https://fred.stlouisfed.org/series/PCU33312033312011

A Derivations of the impact of parameters on the wage peak

Comparative statics of the automation turning point α_I^*

We write the turning-point share as

$$\alpha_I^* \;=\; \big[A\,B\big]^{E}, \quad A=\frac{\rho_I\,\tau}{\theta_I(1-\rho_I)(1-\rho)(1-\tau)}, \quad B=\Big(\tfrac{\bar{P}}{K_I^{\theta_I}}\Big)^{\rho-1}, \quad E=\frac{\rho_I}{\theta_I(1-\rho_I)(\rho-1)},$$

under the maintained parameter domains $0<\rho_I<1,\ \rho<1,\ 0<\tau<1,\ \theta_I>0$. Let $\Lambda\equiv\ln A+\ln B$.

1. Log-linearisation

$$\ln\alpha_I^* \;=\; E\,\Lambda, \qquad \Longrightarrow \qquad \frac{d\alpha_I^*}{dx}=\alpha_I^*\left[\frac{dE}{dx}\,\Lambda+E\,\frac{d\Lambda}{dx}\right]. \tag{22}$$

2. Building blocks

$$E = \frac{\rho_I}{\theta_I(1-\rho_I)(\rho-1)} \quad \Longrightarrow \quad \frac{dE}{dx} = \begin{cases} -\dfrac{E}{\theta_I}, & x = \theta_I, \\ \dfrac{1}{\theta_I(1-\rho_I)^2(\rho-1)}, & x = \rho_I, \\ -\dfrac{E}{\rho-1}, & x = \rho, \\ 0, & x = \tau; \end{cases} \tag{23}$$

$$\ln A = \ln\rho_I + \ln\tau - \ln\theta_I - \ln(1-\rho_I) - \ln(1-\rho) - \ln(1-\tau)$$

$$\Longrightarrow \quad \frac{d}{dx}(\ln A) = \begin{cases} -\dfrac{1}{\theta_I}, & x = \theta_I, \\ \dfrac{1}{\rho_I} + \dfrac{1}{1-\rho_I}, & x = \rho_I, \\ \dfrac{1}{1-\rho}, & x = \rho, \\ \dfrac{1}{\tau} + \dfrac{1}{1-\tau}, & x = \tau; \end{cases} \tag{24}$$

$$\ln B = (\rho-1)\big(\ln\bar{P} - \theta_I \ln K_I\big) \quad \Longrightarrow \quad \frac{d}{dx}(\ln B) = \begin{cases} -(\rho-1)\ln K_I, & x = \theta_I, \\ 0, & x = \rho_I, \\ \ln\bar{P} - \theta_I \ln K_I, & x = \rho, \\ 0, & x = \tau. \end{cases} \tag{25}$$

Combine (24) and (25): $\dfrac{d\Lambda}{dx} = \dfrac{d}{dx}(\ln A) + \dfrac{d}{dx}(\ln B)$.

3. Comparative-static derivatives and signs

$$\frac{\partial\alpha_I^*}{\partial\theta_I} = \alpha_I^*\, E\Big[-\frac{\Lambda+1}{\theta_I} - (\rho-1)\ln K_I\Big] \quad \begin{cases} <0, & K_I < K_I^\dagger \\ =0, & K_I = K_I^\dagger \\ >0, & K_I > K_I^\dagger \end{cases} \tag{$D_{\theta I}$}$$

$$\frac{\partial\alpha_I^*}{\partial\rho_I} = \alpha_I^*\Big[\frac{\Lambda}{\theta_I(1-\rho_I)^2(\rho-1)} + E\Big(\frac{1}{\rho_I} + \frac{1}{1-\rho_I}\Big)\Big] \;<\; 0 \tag{D_{ρ_I}}$$

$$\frac{\partial\alpha_I^*}{\partial\rho} = \alpha_I^*\Big[-\frac{E\,\Lambda}{\rho-1} + E\Big(\frac{1}{1-\rho} + \ln\bar{P} - \theta_I\ln K_I\Big)\Big] \;<\; 0 \tag{D_{ρ}}$$

$$\frac{\partial\alpha_I^*}{\partial\tau} = \alpha_I^*\, E\Big(\frac{1}{\tau} + \frac{1}{1-\tau}\Big) \;<\; 0 \tag{D_{τ}}$$

AI-capital threshold: $K_I^\dagger = \exp\big[-1/((\rho-1)\theta_I)\big] > 1$ delineates the sign change in $\partial\alpha_I^*/\partial\theta_I$.

4. Economic interpretation (sign logic)

- **Scale in I (θ_I).** Scarce AI ($K_I < K_I^\dagger$) $\Rightarrow$ higher θ_I moves the wage peak earlier; abundant AI ($K_I > K_I^\dagger$) $\Rightarrow$ delays it.
- **Task substitutability in I (ρ_I).** Larger ρ_I weakens within-block scale gains $\Rightarrow$ peak arrives earlier.
- **Cross-block substitutability (ρ).** Making P and I easier to substitute ($\rho \uparrow$) reduces the marginal value of extra I $\Rightarrow$ earlier peak.
- **Physical weight (τ).** A higher physical share lowers the leverage of I in output $\Rightarrow$ earlier peak.

5. Sign summary table

Parameter $\uparrow$	Sign of derivative	Qualitative effect on α_I^*
θ_I (if $K_I < K_I^\dagger$)	< 0	Peak earlier (scarce AI)
θ_I (if $K_I > K_I^\dagger$)	> 0	Peak later (abundant AI)
ρ_I	< 0	Peak earlier
ρ	< 0	Peak earlier
τ	< 0	Peak earlier

B Share of earnings in manual occupations in 2023

This appendix estimates the share of total U.S. wage earnings attributable to manual occupations in 2023.
Manual occupations are defined as major occupational categories involving physical labor or skilled trades. These include:

- Food Preparation and Serving
- Building and Grounds Cleaning and Maintenance
- Personal Care and Service
- Farming, Fishing, and Forestry
- Construction and Extraction
- Installation, Maintenance and Repair
- Production
- Transportation and Material Moving

The categorization follows major SOC groupings from the U.S. Bureau of Labor Statistics Occupational Employment and Wage Statistics dataset (U.S. Bureau of Labor Statistics, 2023). BLS documentation explains that classification depends on the *nature of duties*, particularly whether work is physical, hands-on, and involves machinery (U.S. Bureau of Labor Statistics, 2001).

While this estimate uses SOC-based groupings, the approach aligns conceptually with the task-based framework of Autor et al. (2003), who classify work based on whether tasks are routine or non-routine, cognitive or manual. The manual categories used here largely reflect their definition of routine and non-routine manual labor, though we do not directly use task intensity scores from O*NET or other micro sources.

B.1 Wage Bill Calculation (2023)

Table 4: Employment and Wages by Manual Occupation (May 2023)

SOC	Description	Employment	Mean Wage ($)	Wage Bill (Billion $)
00-0000	All Occupations	151,853,870	65,470	9,941.87
35-0000	Food Prep & Serving	13,247,870	34,490	456.92
37-0000	Building & Grounds Cleaning & Maintenance	4,429,070	38,320	169.72
39-0000	Personal Care & Service	3,040,630	38,430	116.85
45-0000	Farming, Fishing & Forestry	432,200	39,970	17.28
47-0000	Construction & Extraction	6,225,630	61,500	382.88
49-0000	Installation, Maintenance & Repair	5,989,460	58,500	350.38
51-0000	Production	8,770,170	47,620	417.64
53-0000	Transportation & Material Moving	13,752,760	46,690	642.12

B.2 Manual Occupation Calculations

We compute key ratios and shares using the following calculations:

- **Ratio of Manual Employment to Cognitive Employment:**

$$\text{Manual Employment} = 55{,}887{,}790,$$

$$\text{Cognitive Employment} = 151{,}853{,}870 - 55{,}887{,}790 = 95{,}966{,}080$$

$$\text{Ratio} = \frac{55{,}887{,}790}{95{,}966{,}080} \approx 0.582$$

  This means that for every 10 cognitive workers, there are approximately 5.8 manual workers in the U.S. labor market in 2023.

- **Share of Manual Employment in Total Employment:**

$$\frac{55{,}887{,}790}{151{,}853{,}870} \approx 0.368 \quad \text{or} \quad 36.8\%$$

- **Share of Manual Occupations in Total Wage Bill:**

$$\text{Manual Wage Bill} \approx \$2{,}554 \text{ B}, \quad \text{Total Wage Bill} \approx \$9{,}942 \text{ B}$$

$$\frac{2{,}554}{9{,}942} \approx 0.2569 \quad \text{or} \quad 25.69\%$$

B.3 Final estimate

Manual occupations account for approximately 25.69% of total U.S. wage earnings in 2023. This estimate is based on aggregating major occupational groups that involve primarily physical labor and hands-on tasks, as defined by the U.S. Bureau of Labor Statistics. These include food service, construction, maintenance, production, and transportation-related work, among others.

While this approach does not apply the full task-based intensity metrics developed by Autor et al. (2003), it provides a replicable, transparent method using standard SOC classifications and wage/employment statistics. The categorization aligns conceptually with the blue-collar definition in the Occupational Classification System Manual (OCSM) (U.S. Bureau of Labor Statistics, 2001), which emphasizes physical effort and machine usage as key dimensions of manual labor.