Discussion Paper Series – CRC TR 224

Discussion Paper No. 804
Project C 05

Creative Destruction With AI Versus Ai's Destruction of Creativity

Christian Bayer[1]
Petr Sedláček[2]

November 2026

[1]University of Bonn, Department of Economics, Adenauerallee 24–42, 53113 Bonn, Germany; CASFI, CEPR, IZA, and CESifo. Email: christian.bayer@uni-bonn.de.
[2]University of New South Wales, Business School, Kensington, NSW 2052, Australia; CASFI and CEPR. Email: p.sedlacek@unsw.edu.au.

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Creative Destruction with AI

versus

AI's Destruction of Creativity*

Christian Bayer† Petr Sedláček‡§

November 11, 2026

Abstract

Debates about AI have largely focused on its effects on productivity within existing firms or tasks. We argue that a key question for long-run welfare is whether AI will itself generate new ideas and products. Using a model of firm entry and selection where product development is uncertain and exclusively human-led, we show that AI automation may reduce welfare, even if it also raises productivity. This happens because automation erodes private returns that help justify investment into new ideas. Therefore, higher AI-induced productivity need not increase welfare if it undermines Schumpeterian incentives for creativity.

JEL codes: E23, E22, L11, D61, O31, O33.

Keywords: Firm entry; product variety; artificial intelligence; automation; augmentation; innovation; markup dispersion.

*Sedláček gratefully acknowledges the financial support of the Australian Research Council (grant number FT230100545). Christian Bayer gratefully acknowledges support from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC2126/2 – 390838866, KFG-50 (CASFI), and through CRC TR 224 (Project C05).

†University of Bonn, Department of Economics, Adenauerallee 24–42, 53113 Bonn, Germany; *CASFI, CEPR, IZA, and CESifo.* email: christian.bayer@uni-bonn.de.

‡University of New South Wales, Business School, Kensington, NSW 2052, Australia; *CASFI and CEPR.* Email: p.sedlacek@unsw.edu.au,

§This project was inspired by discussions at CASFI's "Schumpeter@100" conference 2026. Christian Bayer thanks UNSW for its hospitality as part of this was written while he was visiting UNSW. During the preparation of this work the authors used ChatGPT-6 in order to improve readability and language as well as for checking and improving algebra. The authors reviewed and edited the content afterwards as needed and take full responsibility. All remaining errors are our own.

1 Introduction

Would Intel have developed its first microprocessor when it did without the prospect of making money selling the chips for calculators? Would Dvořák have studied music without the prospect of earning a living in cafés and orchestras?[1] Exceptional innovations depend partly on the prospect of ordinary commercial success. So far, people have developed the ideas behind new products. Whether AI will become a source of genuinely new ideas is an open and understudied question. In this paper, we argue that the answer to this question has profound implications for society's long-run welfare.

Our analysis rests on two key features. First, the development of new products or ideas is done by people. Second, success of such new products and ideas is uncertain, yet investment into them happens upfront. Firms develop products, and individuals acquire skills, before knowing how successful they will be. Importantly, the prospect of a modest commercial return or a decent living without becoming a "superstar" is an integral part of the expected return to these investments. It can make them worthwhile even when exceptional success is unlikely.[2]

We argue that this nexus is important when thinking about the effects of AI on productivity and growth. If AI competes away ordinary returns, it may discourage entry and reduce the number of exceptional products ultimately developed. Yet, AI may also increase the rewards to exceptional success and thereby encourage entry. The history of recorded music illustrates both forces. Recording technology displaced live performers (American Federation of Musicians, n.d.) while allowing successful artists to reach much larger audiences and earn much higher returns (Rosen, 1981).[3] Larger rewards at the top can sustain investment in musical careers and the development of new genres even as opportunities for the average musician to earn a modest living diminish.

We study these opposing forces in a model of firm entry and selection building on Melitz (2003). This framework puts the Schumpeterian idea of productivity being driven by the development and entry of products front and center. Firms (or equivalently households) decide to develop new products, not knowing how productive they will be in producing them. Fixed costs of development and overhead costs of operation shape entry and exit decisions. Through *augmentation*, AI can increase the rewards to successful human innovation. Through *automation*, AI competes away developers' returns, regardless of whether it replaces them entirely as producers or merely erodes their market share. Be-

[1] For the historical references, see Intel Corporation (n.d.) and Šupka (n.d.).

[2] Baumol and Baumol (1994) argue that the abundance of court appointments in eighteenth-century Central Europe attracted people into music and provided opportunities for their talent to emerge.

[3] In his successful Viennese years, Mozart, arguably the musical superstar of his age, earned roughly ten times the salary of a court orchestral musician (Waldoff, 2018). Taylor Swift's estimated earnings in 2025 were nearly three thousand times the average American wage (Di Licosa and Craig, 2025; Social Security Administration, n.d.), using Forbes' earnings estimate and the Social Security Administration's average wage data. Although the benchmarks differ, the comparison illustrates the much larger rewards to musical stardom in the era of mass distribution.

cause firms invest before learning their productivity, entry depends on returns across the entire distribution of possible outcomes. The balance of these forces determines whether AI encourages or discourages entry and, consequently, how many exceptional products are developed. This response shapes the long-run economic gains from AI.

We find that AI augmentation leads to aggregate productivity benefits larger than the direct productivity gains at the individual firm level. As the economy as a whole becomes more productive, the relative costs of further innovation decline, and the number of available products increases. This creates a positive AI-productivity multiplier. This effect becomes stronger if productivity gains from augmentation concentrate at the top.

Automation, by contrast, can have devastating effects on the economy's productivity in the long run. Recall that, because productivity is uncertain at the time of development, some ideas turn out to be commercially only ordinary successes (or even failures), while occasionally ideas rise to exceptional commercial success. In this environment, we start off with a seemingly benevolent scenario in which AI automates only failed ideas, making them equally productive as originally viable human goods and charging the same markups. [4] We find that even though this scenario only adds "free" ideas, the aggregate productivity remains, at best, unchanged. This happens because such free ideas obtain a market share. Therefore, returns to human innovators of initially commercially viable goods decline. In fact, they decline to the level where, despite the free product-designs, the total number of products in the economy is unchanged.

Moreover, if there is fiercer competition leading to a decline in markups among AI designs, the resulting contraction in development more than offsets the direct productivity gain. In such a case, gross output per worker and consumption fall. Interestingly, this negative result worsens as the AI techniques become more productive because they capture larger market shares.

This negative effect becomes even worse if AI automates "ordinary" human innovations which lead to products with small markets and profits. In such a case, AI automation competes away profits from these ordinary outcomes and further weakens incentives to invest into development. In this case, we find that steady state welfare (consumption) falls if ordinary success makes up a sufficiently high share of expected returns to idea development. In fact, in the presence of fierce competition (and resulting markup dispersion), we show that welfare always declines.

While we obtain the above results for a setup in which ideas come with only binary outcomes – exceptional goods vs ordinary goods (or failures) – we also show that our insights carry over to more general settings. In particular, we show that AI automation typically decreases entry and thereby consumption with arbitrary productivity distributions. Finally, we close our analysis by discussing transitions between economies without and with AI and economies where AI both augments and automates production.

[4] As a (non-commercial) example, think about OpenAI solving a Millennium Problem in mathematics, an idea its originator could not solve.

The remainder of the paper is organized as follows. After a literature review, Section 2 introduces the model of product development and firm selection. Section 3 studies augmentation, distinguishing uniform productivity gains from gains concentrated among more or less productive firms. Section 4 turns to automation. We first use three examples to separate the effects of making unsuccessful ideas viable, displacing ordinary commercial success, and introducing markup differences. We then establish the comparison with no AI for general human and conditional AI productivity distributions and positive operating costs. Section 5 summarizes robustness and the transition dynamics of the model. Section 6 discusses the implications of the results. An appendix follows.

Related literature. Our paper relates to several strands of the literature. We connect to papers studying the aggregate effects of changes in firm-level productivity and endogenous product development (see e.g. Baldwin and Robert-Nicoud, 2008; Atkeson and Burstein, 2010; Aghion et al., 2023). The latter papers show that the positive welfare effects stemming from trade liberalizations, more stringent competition among innovators or lower operating costs can largely be undone by the associated decline in product innovation. While our analysis shares the distinction between the efficiency of producing existing products and the incentives to develop new ones, we examine this trade-off in an environment where goods need not be supplied by the original developer, introducing a new channel through which returns to product development can be eroded.

Relatedly, Acemoglu and Restrepo (2018) study the labor market effects of automation alongside the creation of new tasks in which labor has a comparative advantage. Instead, our framework focuses on how AI changes the returns to developing new ideas when development attempts cannot be directed towards markets immune to AI automation.

Our focus on development incentives also relates to papers studying the role of uncertainty for investment (see e.g. Bonfiglioli et al., 2018; Aguiar and Waldfogel, 2018). These papers stress that the distribution of rewards, and uncertainty around them, helps shape incentives to develop (particular) goods. In these frameworks, lower entry costs can, therefore, generate benefits beyond the addition of marginal products. Our analysis, instead, focuses on highlighting that eroding rewards for ordinary goods can have negative spillover effects on incentives to develop exceptional goods.

This paper also connects to studies of production efficiency and efficiency of entry and exit among heterogeneous firms (see e.g. Dhingra and Morrow, 2019; Edmond et al., 2023). In contrast, in our setting AI changes not only the allocation of production across existing suppliers, but also the returns accruing to developers. Importantly, while standard policy intervention can remove markup dispersion and restore efficient allocations across producers, it does not protect developer returns from erosion by AI.

Recently, an emerging literature studies how AI may affect production and the incentives to innovate. Aghion et al. (2019) analyze how AI can discourage human research

by accelerating creative destruction and discuss the consequences of easier imitation. Relatedly, Bazzichi et al. (2026) study how AI affects the choice between incremental and radical innovation when it facilitates research but also shortens the period over which innovators earn profits. Korinek et al. (2026) develop macroeconomic scenarios for the effects of AI automation and augmentation on output, wages, and employment. In contrast to these studies, we focus on analyzing how the AI-induced erosion of ordinary commercial returns can reduce the number of exceptional discoveries.

This also links our analysis to the longstanding literature on the divergence between private and social returns to innovation (Arrow, 1962) and the role of patent protection in sustaining innovation incentives (Nordhaus, 1969; Gilbert and Shapiro, 1990; Klemperer, 1990; Gallini, 1992). In our model, AI can increase the social value of a developed idea by making an otherwise unviable product feasible or by lowering its production cost. At the same time, automation can decrease or eliminate the developer's private return when AI suppliers take over production without compensating the developer. In this case, AI increases the wedge between private and social returns to innovation.

Finally, while not modeled explicitly in our framework, we believe that a closely related argument concerns the acquisition and transmission of expertise. For instance, Garicano and Rayo (2025) show how automation threatens career viability when routine work helps sustain training, Ide (2025) studies how automation changes the transmission of knowledge from experienced workers to novices, Acemoglu et al. (2026) show how AI can weaken human learning incentives and thereby erode collective knowledge, and Asriyan et al. (2026) show that automating junior tasks can lower long-run output by weakening incentives to enter careers in which experience raises productivity.[5] In our framework, instead, ordinary and exceptional returns are different realizations of the same investment. Therefore, our mechanism does not require a training relationship, a financing constraint, or a learning externality. Instead, the erosion of ordinary returns makes overall investment (including eventually exceptional realizations) less attractive.

2 A Tractable Model of Firm Selection

In this section, we introduce our baseline economy. Building on Melitz (2003), in our framework firms must first pay a cost to develop a product. After development, firms learn their production efficiency and, since production entails a fixed cost, they decide whether or not to produce. In the next section, we introduce AI into this baseline economy and trace out its impact on aggregate productivity, firm entry and exit.

[5] Our analysis also relates to work on AI and the supply of creative content (see e.g. Katsamakas and Sanchez-Cartas, 2024).

2.1 Household

A representative household owns all firms and supplies L units of labor each period. The wage is normalized to one. We focus on and compare stationary equilibria. Household preferences are $\sum_{t=0}^{\infty} \beta^t C_t$, where $\beta \in (0,1)$ is the discount factor and C_t is consumption.

The composite good. Let Ω denote the set of supplied products (ω) and M its mass. Gross output (the composite good) is a CES aggregate of these products:

$$Y := \left[\int_{\Omega} y(\omega)^{(\sigma-1)/\sigma}\, d\omega\right]^{\sigma/(\sigma-1)}, \qquad P := \left[\int_{\Omega} p(\omega)^{1-\sigma}\, d\omega\right]^{1/(1-\sigma)}, \tag{1}$$

where P is the price of the composite good, $p(\omega)$ is the price of an individual product and σ is the elasticity of substitution between goods. In what follows, we assume $\sigma > 2$, which, as will become clear, implies diminishing returns to product variety and ensures that the economy's scale is well behaved when the economy uses its own output to develop products.

Demand for individual goods. Minimizing the expenditure needed to assemble gross output, $\min_{\{y(\omega)\geq 0\}_{\omega\in\Omega}} \int_{\Omega} p(\omega)y(\omega)\, d\omega$ subject to $\left[\int_{\Omega} y(\omega)^{(\sigma-1)/\sigma}\, d\omega\right]^{\sigma/(\sigma-1)} \geq Y$, yields familiar demand for individual goods:

$$y(\omega) = Y\left(\frac{p(\omega)}{P}\right)^{-\sigma}. \tag{2}$$

2.2 Firms

In order to produce a good, potential entrants must first pay a product *development cost* $\kappa_d > 0$. In addition, producing a developed design requires payment of a fixed *operating cost* $\kappa_o \geq 0$ every period. Both costs are in units of the composite good. Finally, at the end of each period, goods exogenously fall out of fashion with probability $\delta \in (0,1)$, losing all value.

Producing firms and firm exit. A good ω is produced by an individual firm according to $y(\omega) = z(\omega)\ell(\omega)$, where $z(\omega)$ is firm-level productivity (assumed constant over time) and $\ell(\omega)$ is production labor.[6] Given that individual productivity is constant over a firm's lifetime, operating profits $\pi(z) = p(z)y(z) - \ell(z) - P\kappa_o$ are static.

We assume that firms which have opted to stay in the economy optimally choose prices and labor demand to maximize firm value (the discounted stream of profits), $v(z) = \max_{p,\ell} \pi(z)/\rho$, where $\rho = 1 - \beta(1-\delta)$ is the effective discount rate. In this setting, firms

[6]Productivity need not necessarily relate to quantity only, but could be thought of, alternatively, as product quality, q, that shifts demand, with $Y := \left[\int_{\Omega} q(\omega)\ell(\omega)^{(\sigma-1)/\sigma}\, d\omega\right]^{\sigma/(\sigma-1)}$, $z(\omega) := q(\omega)^{\sigma/(\sigma-1)}$.

that decide to remain in business optimally set prices as a constant markup over marginal cost, $p(z) = \mu/z$, where $\mu = \sigma/(\sigma-1)$, and employ workers in order to satisfy demand for their product $\ell(z) = y(z)/z$.

In addition to firms shutting down when their product exogenously falls out of fashion, businesses may also endogenously choose to exit. This happens when firm-level productivity is not sufficiently high to cover the fixed cost of operation and firm value falls below zero. Therefore, the productivity of the marginal firm, z^*, is implicitly defined by

$$\pi(z^*) = \mu^{1-\sigma} P^\sigma Y \frac{\mu-1}{\mu} (z^*)^{\sigma-1} - P\kappa_o = 0. \tag{3}$$

For future reference, it will be convenient to define the weighted productivity average of producing firms by

$$\tilde{z} := \left[\frac{1}{M}\int_\Omega z(\omega)^{\sigma-1} d\omega\right]^{1/(\sigma-1)}. \tag{4}$$

With this notation, observe that expected profits equal profits evaluated at average productivity:

$$\frac{1}{M}\int_\Omega \pi[z(\omega)] d\omega = \pi(\tilde{z}). \tag{5}$$

Firm entry and the development of goods. As highlighted above, a potential entrant must first pay a cost, $\kappa_d > 0$, to develop a product. Upon development, firms draw (constant) productivity levels from a distribution G and decide whether to produce or rather exit endogenously. Defining $s^* := 1 - G(z^*)$ as the fraction of developed products which are eventually supplied to the market, we can write the familiar free entry condition:

$$\rho P \kappa_d = \int_{z^*}^\infty \pi(z) dG(z) = s^* \pi(\tilde{z}). \tag{6}$$

The free entry condition (6) implicitly defines total development attempts, I. Using K to denote the stock of developed products still in fashion (including designs that are not produced), the mass of supplied products is given by $M \equiv s^* K$ and stationary replacement investment satisfies $I = \delta K$.

2.3 Aggregation and Market Clearing

Given that all labor is used for production and all costs are paid in units of the composite good, labor market clearing and the resource constraint are given by:

$$\int_\Omega \ell(\omega)\, d\omega = L, \qquad Y \equiv C + \kappa_d I + \kappa_o M. \tag{7}$$

Notice that, in this framework, development of products competes with consumption for goods. As a result, an increase in development activity, ceteris paribus, decreases avail-

able consumption. Yet, by adding varieties, it also increases productivity and, thereby, resources available for consumption.

The optimal price setting behavior of firms gives rise to an aggregate representation of our economy. Using the demand schedule of all firms and their optimal price setting, we obtain for total labor demand (and labor market clearing)

$$L = \int_\Omega \ell(\omega)d\omega = Y\int_\Omega \left(\frac{p(\omega)}{P}\right)^{-\sigma} z(\omega)^{-1}d\omega = Y\left[\int_\Omega z(\omega)^{\sigma-1}d\omega\right]^{1/(1-\sigma)}$$

This means we can write the aggregate production function Y for the economy and the real wage P^{-1} as

$$Y = M^{1/(\sigma-1)}\tilde{z}L, \qquad P^{-1} = \frac{M^{1/(\sigma-1)}\tilde{z}}{\mu}. \tag{8}$$

Since the average (not exiting) firm produces at productivity $\tilde{z}$, independent of the number of varieties in the economy, the expected profits (in goods) before operating expenses of a successful developer are given by (left-hand side)

$$\frac{\mu-1}{\mu}Y/M = \frac{\rho\kappa_d}{s^*} + \kappa_o, \tag{9}$$

and equal the cost of (successful) development plus operating costs (right-hand side).[7] Therefore, in our baseline model the equilibrium number of varieties is given by

$$M_b = \left[\frac{\sigma^{-1}\tilde{z}L}{\rho\kappa_d/s_b^* + \kappa_o}\right]^{1+1/(\sigma-2)}. \tag{10}$$

As a result, (holding z^* and thus s^* fixed) an increase in productivity to $\tilde{z}'$ increases the number of varieties and, therefore, aggregate labor productivity increases by more than factor $\tilde{z}'/\tilde{z}$, see (8).

Finally, consumption is given by the resource constraint and is proportional to output in the patient limit ($\rho = \delta$) or when operating costs, κ_o, are zero (fixing $s^* = 1 - G(0)$):

$$\begin{aligned} C_b &= Y_b - \kappa_o M_b - \delta\frac{M_b\kappa_d}{s_b^*} = M_b\left[(\sigma-1)\kappa_o + \frac{(\sigma\rho-\delta)\kappa_d}{s_b^*}\right] \\ &= \frac{\sigma-1}{\sigma}Y_b + (\rho-\delta)M_b\frac{\kappa_d}{s_b^*} = \left(\frac{\sigma-1}{\sigma} + \frac{\rho-\delta}{\rho\sigma}\right)Y_b - \frac{\rho-\delta}{\rho}\kappa_o M_b. \end{aligned} \tag{11}$$

3 AI Augmentation

In this section, we first consider the case where AI improves productivity by "augmenting" production by existing firms. In this setting, the developer retains its product and simply uses an improved technology.

[7] The markup implies that the labor resource cost of product development is μ^{-1} times its private cost. With the wage normalized to one, producing the κ_d development goods requires $P\kappa_d/\mu$ units of labor, whereas the developer pays $P\kappa_d$.

3.1 Setup

AI augmentation in production. As in the baseline economy without AI, we assume that upon developing a product, firms draw (unaugmented) production efficiency u from distribution G. In the presence of augmentative AI, however, effective firm-level productivity is given by $z = \alpha(u)u$, where $\alpha(u) \geq 1$. In addition, we assume $\alpha'(u)u + \alpha(u) > 0$, so augmentation preserves the original productivity ordering, and that the relevant moment of effective productivity is finite. The remainder of the environment is unchanged relative to the baseline economy without AI.

In what follows, we will use a convenient reformulation of the above setting whereby AI transforms the distribution of (unaugmented) productivity, G, into the distribution of effective productivity, $G_\alpha(\alpha(u)u) \equiv G(u)$. We use subscript α to denote the AI-augmented economy and subscript b to denote the "baseline" economy without AI.

Equilibrium with AI augmentation. We now show that all equilibrium conditions in the AI-augmented economy are the same as in the benchmark without AI when evaluated under the distribution of effective productivity, G_α.

In particular, effective productivity of the marginal firm is implicitly defined by $\pi(z^*_\alpha) = 0$, see (3). Next, the fraction of supplied products and the average productivity index are given by:

$$s^*_\alpha := 1 - G_\alpha(z^*_\alpha), \tag{12}$$

$$\widetilde{z}_\alpha := \left[\frac{1}{s^*_\alpha}\int_{z^*_\alpha}^{\infty} z^{\sigma-1}\, dG_\alpha(z)\right]^{1/(\sigma-1)}. \tag{13}$$

Finally, free entry and labor market clearing have the same forms as in the benchmark without AI, evaluated at the transformed distribution, G_α:

$$\rho P_\alpha \kappa_d = s^*_\alpha \pi(\widetilde{z}_\alpha), \qquad L = M_\alpha^{\frac{1}{1-\sigma}} \frac{Y_\alpha}{\widetilde{z}_\alpha}. \tag{14}$$

3.2 Impact of AI Augmentation

In the remainder of this section, we provide theoretical results regarding the impact of AI on aggregate productivity, product development, and firm exit (selection). We start with uniform augmentation, $\alpha(u) = A > 1$, which raises every firm's productivity by the same proportion. We then move on to non-uniform augmentation, where AI affects firms differently depending on their (unaugmented) productivity.

Uniform augmentation. The following proposition describes the impact of AI on the economy under uniform augmentation.

PROPOSITION **1** (Uniform augmentation)
Relative to the baseline economy, uniform AI augmentation ($\alpha(u) = A > 1$) leads to:

i) a higher effective productivity cutoff and average supplier productivity, but unchanged share of supplied products

$$z_\alpha^* = Az_b^*, \qquad \widetilde{z}_\alpha = A\widetilde{z}_b, \qquad \frac{s_\alpha^*}{s_b^*} = 1,$$

ii) higher product variety, development investment, real wage, gross output and consumption

$$\frac{M_\alpha}{M_b} = \frac{I_\alpha}{I_b} = \frac{P_b}{P_\alpha} = \frac{Y_\alpha}{Y_b} = \frac{C_\alpha}{C_b} = A^{\frac{\sigma-1}{\sigma-2}}.$$

Proof. Since under uniform augmentation $\alpha(u) = A$, we can write $G_\alpha(z) = G(z/A)$ and, therefore, the free-entry cutoff condition is fulfilled for $z_\alpha^* = Az_b^*$. From (12) and (13) it then follows that $s_\alpha^* = s_b^*$ and $\tilde{z}_\alpha = A\tilde{z}_b$. Equation (10) then gives $M_\alpha/M_b = A^{(\sigma-1)/(\sigma-2)}$ and the same ratio also applies for development investment since $I_\alpha = \delta K_\alpha = \delta M_\alpha/s_\alpha^*$. Equations (9) and (11) show that, for an unchanged fraction of supplied products, gross output and consumption are proportional to the mass of supplied products. Finally, (8) shows that the aggregate price index (inverse of the real wage) is proportional to $1/(M_\alpha^{1/(\sigma-1)}\widetilde{z}_\alpha)$ and, therefore, $P_b/P_\alpha = A^{(\sigma-1)/(\sigma-2)}$. □

Intuitively, Proposition 1 shows that since uniform augmentation raises effective productivity of all firms, it increases productive capacity and lowers the price of goods. Relative to the cost of labor, this lowers the labor costs of product development, $P\kappa_d$, and operation, $P\kappa_o$,. In turn, this leads to an increase in the mass of developed products, even though the fraction of supplied products (selection) remains unchanged.

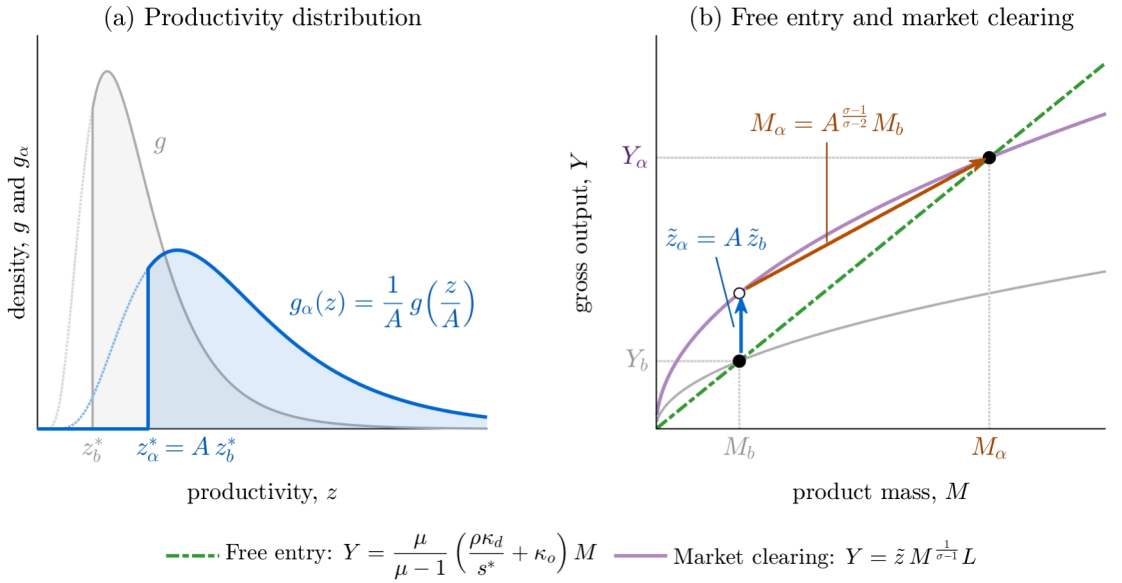
Figure 1 illustrates this amplification. Panel (a) shows that uniform augmentation raises productivity and the operating cutoff proportionally, leaving the fraction of supplied products unchanged. Panel (b) separates the increase in output at a given number of products from the additional increase generated by entry.

Put together, the consumption gain exceeds the direct productivity improvement. Specifically, consumption increases by a factor of A due to higher average productivity. In addition, it rises by an additional factor of $A^{1/(\sigma-2)}$ through greater product variety. Since common markups continue to allocate production efficiently, the share of output available for consumption is unchanged.

Non-uniform augmentation. We now turn to the case where augmentation can differ across firms. Towards this end, define A as the factor by which mean productivity increases:

$$A := \frac{\int_0^\infty z\, dG_\alpha(z)}{\int_0^\infty z\, dG(z)} = \frac{\int_0^\infty \alpha(u) u\, dG(u)}{\int_0^\infty u\, dG(u)} > 1. \tag{15}$$

Figure 1: Uniform augmentation



Notes. The figure illustrates Proposition 1. Panel (a) shows the densities of unaugmented productivity and of effective productivity under uniform augmentation (g and g_α) together with the respective cutoffs. Panel (b) shows free entry (9) and market clearing (8) in terms of product variety, M, and gross output, Y.

This arithmetic mean is taken across all developed designs, before selection. Note that under uniform augmentation, the productivity of all businesses rises by the factor A.

For the cutoff and strict distributional comparisons in the following proposition, assume $\kappa_o > 0$ and a continuous, nondegenerate distribution G of positive productivities.

PROPOSITION **2** (Nonuniform augmentation)
Relative to the baseline economy, non-uniform AI augmentation ($\alpha(u) > 1$) leads to a higher effective productivity cutoff and a higher real wage:

$$\frac{z_\alpha^*}{z_b^*} > 1, \qquad \frac{P_b}{P_\alpha} > 1.$$

Moreover,

(a) If $\alpha'(u) > 0$, the supplied fraction of products weakly decreases, $s_\alpha^* \leq s_b^*$, development investment increases, the number of supplied products can rise or fall, but consumption rises by more than under uniform augmentation with the same mean productivity:

$$\frac{C_\alpha}{C_b} > A^{1+1/(\sigma-2)}.$$

(b) If $\alpha'(u) < 0$, the supplied fraction of products weakly increases, $s_\alpha^* \geq s_b^*$, and the number of supplied products rises. While development investment can rise or fall, consumption rises for β sufficiently close to one, but is lower than under uniform

augmentation with the same mean productivity:

$$\frac{C_\alpha}{C_b} < A^{1+1/(\sigma-2)}.$$

We defer the proof of Proposition 2 to Appendix B and focus only on the intuition. Specifically, the proposition highlights that under non-uniform augmentation, the distribution of productivity gains from AI matters even when augmentation as a whole leads to the same average productivity improvement.

To understand this, consider augmentation that disproportionately rewards high-productivity firms. Because expected profits increase in average productivity and average productivity is the convex $\sigma - 1$-mean of the firm productivity distribution, see (4) and (5), this magnifies expected returns to exceptional outcomes and further incentivizes product development. In this case, development and consumption rise relatively more than under uniform augmentation, even though selection becomes more stringent.

In contrast, if augmentation gains are concentrated among less productive firms, this allows a larger fraction of developed ideas to be supplied. At the same time, however, the distribution of expected profits from product development gets compressed. Therefore, compared with uniform augmentation, the effective productivity of the marginal producer rises by less, and consumption is relatively lower.

4 AI Automation

An alternative, a polar opposite, role for AI is to automate the production of goods or services whose designs have already been developed; replacing the product design and taking over the production activity of the original developer.

AI automation in production. Under automation, AI supplies a competing technology for a developed product. Therefore, it is necessary to distinguish two productivity levels with which a given developed good can be produced.

First, upon paying the development cost κ_d, human-run firms draw productivity z_H from distribution G^H. Second, AI can produce a given good with productivity z_A, drawn from $G^A(\cdot \mid z^H)$. A zero draw (for human-run firms or AI alike) means that the corresponding technology cannot produce. Both draws are observed before the production decision and remain fixed while the product is in fashion. Note that AI can supply any developed idea for which its technology is viable. This potentially includes goods that are not viable for human-run firms.

AI and development. While AI can supply any developed idea, we assume that only humans develop new products (or the idea thereof). As in the baseline economy, product

development is costly. Moreover, developers anticipate competition from AI and expect to receive no royalties if AI has superior technology. In a sense, there is a lack of trademark, copyright, or patent protection against AI—developers cannot exclude close-substitute AI suppliers. As will become clear, (dis)incentives to develop products will be a crucial force driving aggregate outcomes.

Three illustrative examples. Before we turn to the general model, we first consider three illustrative examples to develop intuition. In each of these, we make the following simplifying assumptions. First, we abstract from overhead costs of production, $\kappa_o = 0$. Second, we assume that human productivity can only take two discrete values, low (z_L) and high (z_H), with $0 \leq z_L < z_H$, and normalize $z_H = 1$. The share of high-end products is denoted by $s \in (0, 1)$. Finally, we assume that AI supplies only the $(1 - s)$ share of low-end goods, but does so at a superior productivity, $z_A > z_L$.

As will become clear, the examples highlight two important channels through which AI automation can harm long-run welfare: (i) business stealing and (ii) over-consumption of AI-provided varieties (due to lower markups). In Section 4.4, we show that the results from these illustrative examples carry over to the general model.

4.1 Making Unsuccessful Ideas Useful

In this first example, we consider a case in which AI allows for the production of goods that were previously not viable options for human-run firms.

AI in production. We assume that the $1 - s$ share of low-end goods are not viable for human-run firms, $z_L = 0$. However, AI can produce these goods with productivity $z_A > 0$. In other words, without AI, only a fraction s of developed ideas would be supplied. With AI, every developed idea is supplied. Moreover, note that AI productivity can be lower or higher than that of successful human production, z_H.

In this example, we assume that each ordinary idea is served by a single AI supplier (with a costless exclusive right to serve each failed design) that charges the same monopoly markup as an exceptional good produced by human-run firms. Therefore, as in the baseline economy without AI, prices of all goods are set as constant markups over marginal costs and equal markups imply an efficient allocation across supplied products. The AI supplier is owned by the representative household, too. In Section 4.3, we instead allow for competition among AI suppliers which has important consequences for our aggregate results.

Finally, we assume that the household owns the AI suppliers and receives their profits as dividends. Importantly, however, these profits are not tied to the development of the respective good.

Average productivity. In the baseline economy without AI, average productivity is simply $\widetilde{z}_b = z_H = 1$. With AI, even the fraction $(1-s)$ of previously unviable products is produced (with productivity z_A) and, therefore, average productivity becomes

$$\tilde{z}_A = \left[s + (1-s)z_A^{\sigma-1}\right]^{1/(\sigma-1)}. \tag{16}$$

Labor demand from AI production. Let ℓ_A^A denote production labor per AI-supplied product, ℓ_A^H the labor of human-run firms after AI entry, and ℓ_b^H before. Recalling that $z_H = 1$ in our examples, relative CES demand of AI and human products implies

$$\ell_A^A = z_A^{\sigma-1}\ell_A^H.$$

In what follows, we will use λ to denote the fraction of production labor employed by AI suppliers (which also equals its market share):

$$\lambda = \frac{(1-s)z_A^{\sigma-1}}{s + (1-s)z_A^{\sigma-1}}. \tag{17}$$

Note that in the special case when AI produces goods exactly as efficiently as humans, $z_A = z_H = 1$, the labor share of AI products is simply $1-s$. With $z_A > z_H$, the share of labor in AI goods production increases relative to labor in human production. This happens because the lower price of AI goods expands their demand.

Development of goods. Recall that the free entry condition (6) equalizes expected profits from goods development to the entry cost. In this example, with only two types of productivity levels, one of which is not viable, the free entry condition is given by:

$$s\pi_{A,b}^H = \rho P_{A,b}\kappa_d. \tag{18}$$

A key insight of this example is that even though AI takes over "only" in the market where humans earn zero profits ($\pi^L = 0$), the equilibrium incentives to develop goods will nevertheless change. This happens because AI dilutes the profit share of human-run firms through the reallocation of demand and production labor toward AI products. The following proposition formalizes this insight.

PROPOSITION **3** (Useful failures with common markups)
Consider $\kappa_o = 0$ and AI produces only goods which are unviable for humans and it does so with productivity $z_A > 0$, charging the same markup μ as human suppliers. In this case, relative to the baseline case without AI, automation leads to

i) lower product development

$$\frac{I_A}{I_b} = \frac{K_A}{K_b} = 1 - \lambda.$$

ii) unchanged prices and gross output

$$\frac{P_A}{P_b} = \frac{Y_A}{Y_b} = 1.$$

iii) higher aggregate consumption

$$\frac{C_A}{C_b} = 1 + \lambda \frac{\delta}{\sigma\rho - \delta}.$$

Proof. In the baseline economy without AI, labor market clearing implies $L = sK_b\ell_b^H \Leftrightarrow \ell_b^H = \frac{L}{sK_b}$. With AI automation, labor market clearing is given by

$$L = K_A(s\ell_A^H + (1-s)\ell_A^A) \Leftrightarrow \ell_A^H = \frac{L}{K_A(s + (1-s)z_A^{\sigma-1})}.$$

In turn, using free entry (18) and the fact that $\pi_H = (\mu - 1)\ell_H$, we can write

$$\frac{K_A}{K_b} = (1-\lambda)\frac{P_b}{P_A}. \tag{19}$$

Next, using (8) we can write the aggregate price index in the economy with(out) AI as

$$P_b^{1-\sigma} = sK_b\mu^{1-\sigma}, \qquad P_A^{1-\sigma} = K_A\mu^{1-\sigma}\left[s + (1-s)z_A^{\sigma-1}\right].$$

Taking their ratio and substituting (19) yields

$$\left(\frac{P_b}{P_A}\right)^{\sigma-2} = (1-\lambda)\frac{s + (1-s)z_A^{\sigma-1}}{s} = 1.$$

This, together with the preceding results, proves part (ii). Substituting the above back into (19) proves part (i) of Proposition 3. Equation (8) implies $P_AY_A = P_bY_b = \mu L$. This, together with the above results, proves part (ii) of Proposition 3. Finally, since $Y_A = Y_b$, $I_A = I_b(1-\lambda) < I_b$ and $C = Y - I\kappa_d$, AI automation directly increase consumption by reducing the goods used for idea development ($\lambda\kappa_d I_b$). Specifically, the following

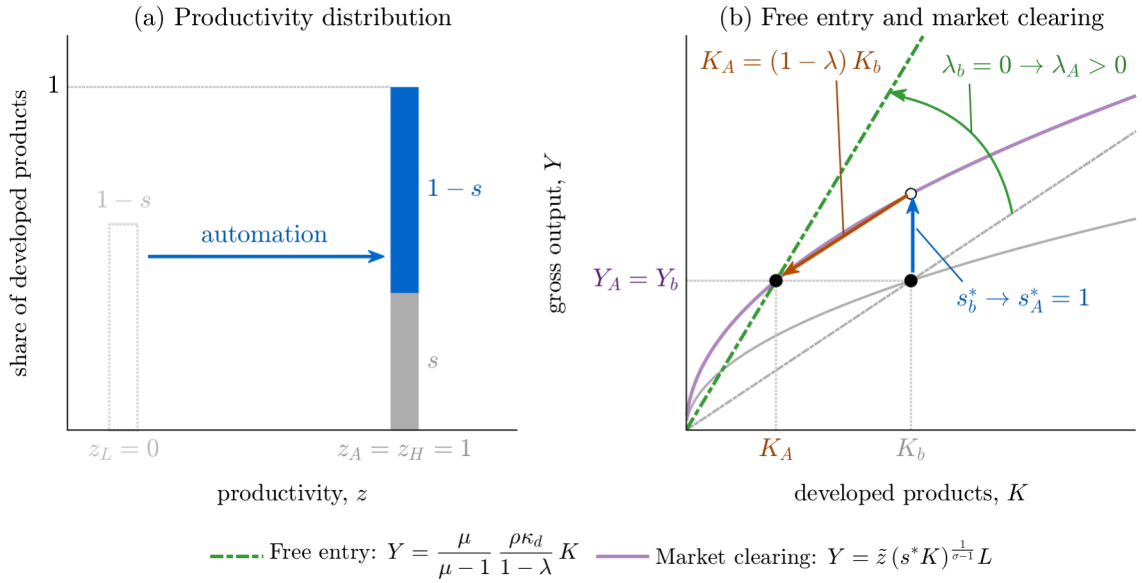
calculation proves part (iii) of Proposition 3:

$$\frac{C_A}{C_b} = 1 + \lambda \frac{\kappa_d I_b}{C_b} = 1 + \lambda \frac{\delta}{\sigma\rho - \delta},$$

where the last equality uses $I_b = \delta K_b$ and (11) to express $C_b = K_b(\sigma\rho - \delta)\kappa_d$. □

Figure 2 illustrates why making unsuccessful ideas useful raises consumption without raising gross output. At a given stock of developed ideas, automation increases output, but it also reduces the share of sales that rewards developers. The stock of developed ideas therefore contracts until gross output returns to its previous level. Maintaining this output requires fewer development goods, leaving more goods available for consumption. The horizontal axis measures developed ideas, K, rather than supplied varieties, M.

Figure 2: Useful failures with common markups



Notes. The figure illustrates Proposition 3. Panel (a) shows the two-point distribution of human productivity, where the share $1-s$ of developed products is not viable for human-run firms and AI supplies these products with productivity z_A. Panel (b) shows free entry (18) and market clearing (8) in terms of developed products, K, and gross output, Y.

Intuitively, this example highlights that consumption gains need not stem from increases in gross output. A useful special case is when AI produces previously unviable products with the same productivity as human-produced goods. In this case, which is visualized in Figure 2, the same number of products is supplied, $M_A/M_b = 1$, with the same production technology and the same labor input per variety, $z_A = z_H = 1$. However, while only a fraction s of developed products were supplied to the market in the baseline economy, all developed goods are produced in the automated one. Therefore, the same mass of produced goods is achieved with fewer developed products, $I_A/I_b = K_A/K_b = s$, see Proposition 3.

This also makes clear that, despite automation not changing gross output relative to the benchmark, consumption rises. This happens because fewer development attempts are needed to sustain the same mass of produced goods since AI can supply previously unviable goods. Note, however, that these consumption gains from automation are limited by the resources that the baseline economy devotes to product development. As AI productivity rises, its market share increases too, see (17). In the extreme case of $z_A \to \infty$, $\lambda \to 1$ and the labor force goes towards a vanishing mass of products. While gross output remains at $Y_A = Y_b$, consumption moves towards output as product development vanishes, $C_A \to Y_b$. Therefore, the relative increase in consumption from AI automation is bounded from above by $Y_b/C_b = \sigma/(\sigma - \delta/\rho)$, see (11). In the limit where households are infinitely patient, $\beta \to 1$, this upper bound is equal to the markup, $Y_b/C_b = \sigma/(\sigma-1) = \mu$.

4.2 Automating Ordinary Commercial Success

In our second example, we allow for all goods to be commercially viable even without AI.

AI in production. We assume that the $1 - s$ share of low-end goods ("ordinary outcomes") can be produced without AI, $0 < z_L < 1$. As in the previous example, AI automation is available only for these products, and each AI product is supplied by a single producer charging the markup $\mu = \sigma/(\sigma-1)$. In addition, we assume $z_A \geq \mu z_L$, so that its price μ/z_A is no higher than the human supplier's marginal cost $1/z_L$. At equality, we assume AI serves the market as a tie-breaking rule.

As before, the high-productivity human provider remains protected from AI and charges the markup μ over marginal costs. While the household receives the AI profits as dividends, the developers of low-end products do not receive royalties.

Average productivity. In contrast to the previous example, AI automation in this case necessarily increases average productivity:

$$\tilde{z}_A = \left[s + (1-s)z_A^{\sigma-1}\right]^{1/(\sigma-1)} \geq \left[s + (1-s)z_L^{\sigma-1}\right]^{1/(\sigma-1)} = \tilde{z}_b. \tag{20}$$

Labor demand from AI production. Let ℓ_b^H denote production labor per high-productivity human product without AI, and let ℓ_A^H denote its counterpart after AI. In addition, let ℓ_A^A denote the labor demand for AI-produced ordinary goods. As in the previous example, we obtain the same relative labor demand and market share expressions:

$$\ell_A^A = z_A^{\sigma-1}\ell_A^H \qquad \text{and} \qquad \lambda = \frac{(1-s)z_A^{\sigma-1}}{s + (1-s)z_A^{\sigma-1}}. \tag{21}$$

Intuitively, higher AI productivity increases its market share. Although producing a unit requires less labor, the lower price expands demand sufficiently to increase production

labor per AI product relative to a human product.

Development of goods. Let us now turn to the free entry condition and the development of goods. Using that $\pi_b^L = (\mu-1)\ell_b^L = (\mu-1)\ell_b^H z_L^{\sigma-1}$ and $\pi_{A,b}^H = (\mu-1)\ell_{A,b}^H$, we can write free entry before and after automation as:

$$\underbrace{(1-s)(\mu-1)z_L^{\sigma-1}\ell_b^H}_{\text{ordinary outcomes}} + \underbrace{s(\mu-1)\ell_b^H}_{\text{exceptional outcomes}} = \rho\kappa_d P_b, \tag{22}$$

$$s\ell_A^H(\mu-1) = \rho\kappa_d P_A, \tag{23}$$

where P_b and P_A are the aggregate price indices before and after automation. As is clear from the above, a key impact of automation is that it removes the ordinary-outcome contribution to expected profits. Consequently, exceptional outcomes must earn more relative to the cost of development.

Equivalently, use $\ell_A^H = \frac{L}{M_A}\frac{1}{s+(1-s)z_A^{\sigma-1}} = \frac{L}{M_A}\tilde{z}_A^{1-\sigma}$ and $\ell_b^H = \frac{L}{M_b}\frac{1}{s+(1-s)z_L^{\sigma-1}} = \frac{L}{M_b}\tilde{z}_b^{1-\sigma}$ for the automation and no-automation cases to express the per-firm high-quality human labor demand as a fraction of total labor. Substituting these expressions and the price index from (8) to into the free entry condition, we obtain

$$(\mu-1)L/M_b = \rho\kappa_d P_b \Leftrightarrow \frac{L\tilde{z}_b}{\sigma\rho\kappa_d} = M_b^{\frac{\sigma-2}{\sigma-1}} \tag{24}$$

$$(1-\lambda)(\mu-1)L/M_A = \rho\kappa_d P_A \Leftrightarrow \frac{Ls\tilde{z}_A^{2-\sigma}}{\sigma\rho\kappa_d} = M_A^{\frac{\sigma-2}{\sigma-1}} \tag{25}$$

Therefore, as the following proposition formalizes, the number of developed ideas must decrease when some products are produced through AI-automation.

Proposition **4** (Ordinary outcomes with a common markup)
Suppose $\kappa_o = 0$ and AI is used to produce only ordinary goods, with productivity $z_A > \mu z_L$, charging the same markup μ as human suppliers. Then automation implies that:

i) the mass of varieties declines in AI productivity, rendering the real wage and output independent of the AI productivity level,

$$M_A = \left(\frac{Ls}{\sigma\rho\kappa_d}\right)^{\frac{\sigma-1}{\sigma-2}}\tilde{z}_A^{1-\sigma}, \qquad P_A^{-1} = \mu^{-1}\left(\frac{Ls}{\sigma\rho\kappa_d}\right)^{\frac{1}{\sigma-2}}, \qquad Y_A = \mu P_A^{-1}L; \tag{26}$$

ii) relative to the baseline case without AI, the real wage, output, the mass of varieties, and development investment are all lower, and

$$\frac{P_b}{P_A} = \frac{Y_A}{Y_b} = \left[\frac{s}{\tilde{z}_b^{\sigma-1}}\right]^{1/(\sigma-2)} < 1, \qquad \frac{M_A}{M_b} = \frac{I_A}{I_b} = (1-\lambda)\frac{P_b}{P_A} < 1; \tag{27}$$

iii) consumption can rise or fall and satisfies

$$C_A = \underbrace{\left[\frac{s}{s+(1-s)z_L^{\sigma-1}}\right]^{\frac{1}{\sigma-2}}}_{\text{uncompensated loss of varieties}} \times \underbrace{\left[1+\lambda\frac{1}{\sigma\rho/\delta-1}\right]}_{\text{saved development costs}} \times C_b. \tag{28}$$

Here, the first term reflects the productivity effect of the varieties lost due to the loss of ordinary returns, and the second term reflects the saved development costs, including for those varieties that are not developed because of higher AI productivity.

Proof. i) Follows directly from (25) solving for M_A and the statements on the real wage P_A^{-1} and output then follow from plugging in this solution in (8).

ii) By the same line of argument, we obtain

$$M_b = \left(\frac{L\tilde{z}_b}{\sigma\rho\kappa_d}\right)^{\frac{\sigma-1}{\sigma-2}}, \qquad P_b^{-1} = \mu^{-1}\left(\frac{L}{\sigma\rho\kappa_d}\right)^{\frac{1}{\sigma-2}}\tilde{z}_b^{\frac{\sigma-1}{\sigma-2}}, \qquad Y_b = \mu P_b^{-1}L. \tag{29}$$

The relative expressions follow directly.

iii) Observe that $\frac{\mu-1}{\mu} = \sigma^{-1}, Y = \mu L P^{-1}$, and rewrite (24) and (25) as

$$(\mu-1)L/P_b = \rho\kappa_d M_b, \Leftrightarrow \frac{Y_b}{\rho\sigma} = \kappa_d M_b \tag{30}$$

$$(1-\lambda)(\mu-1)L/P_A = \rho\kappa_d M_A, \Leftrightarrow (1-\lambda)\frac{Y_A}{\rho\sigma} = \kappa_d M_A. \tag{31}$$

Finally, $C = Y - \delta\kappa_d M$ (because $M = K$) implies

$$\frac{C_A}{Y_A} = 1-(1-\lambda)\frac{\delta}{\sigma\rho}, \qquad \frac{C_b}{Y_b} = 1-\frac{\delta}{\sigma\rho}.$$

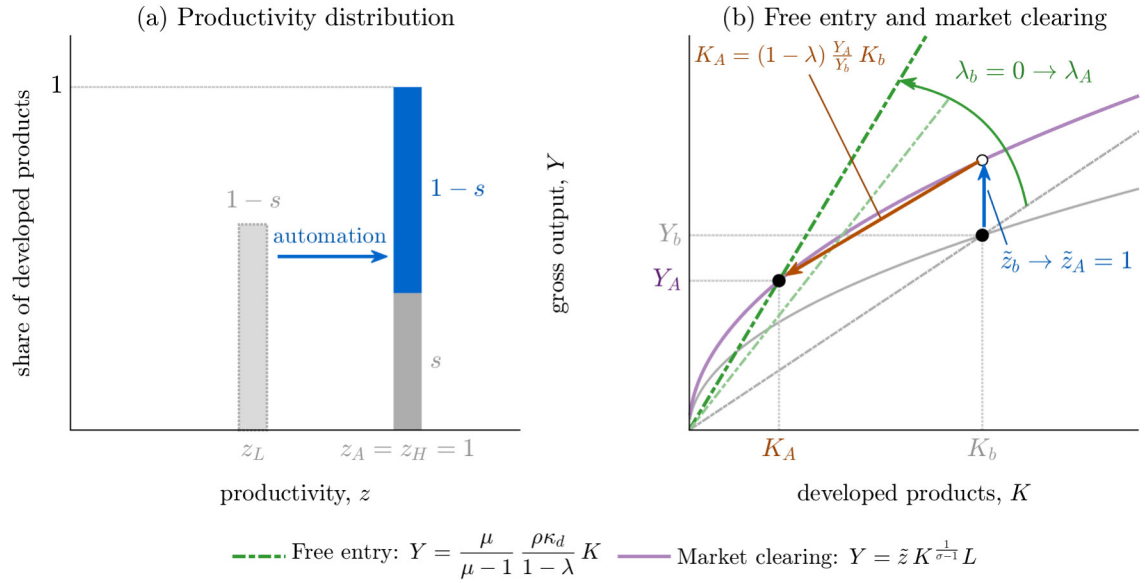
Multiplying their ratio by Y_A/Y_b proves (28).

□

Intuitively, Proposition 4 shows that in the case of automation of previously viable goods, automation reduces incentives to develop products through two channels. This is visualized in Figure 3. First, and as in the previous example of unviable goods, automation dilutes the share of profits going to exceptional goods produced by humans because ordinary products are now produced more efficiently. This effect is indicated by the light-green line in the figure. Indeed, if this were the only effect, the economy would be in the same position as in the previous subsection with $Y_A = Y_b$.

However, a second channel is present when AI automates already viable goods. In particular, AI directly takes over the market share of these ordinary products and reduces the associated expected development profits to zero. This is indicated by a further

Figure 3: Ordinary outcomes with a common markup



Notes. The figure illustrates Proposition 4 for $z_A = z_H = 1$. Panel (a) shows the two-point distribution of human productivity, where AI takes over the share $1-s$ of developed products with ordinary productivity z_L. Panel (b) shows free entry (31) and market clearing (8) in terms of developed products, $K = M$, and gross output, Y. The light green line is free entry with $\lambda' = 1-(\tilde{z}_b/\tilde{z}_A)^{\sigma-1}$, which reflects the dilution of the market share of products with productivity z_H. The remaining tilt reflects the loss of development profits on ordinary products.

tilt of the free entry line, reducing developed products and gross output. At the same time, however, fewer development attempts save the associated costs. Whether aggregate consumption increases or decreases depends on which force dominates. The following corollary formalizes this horse-race.

COROLLARY **1**

A sufficient condition for a decline in consumption after AI automation of ordinary goods is that the share of exceptional goods in expected profits before AI automation, $\frac{s}{s+(1-s)z_L^{\sigma-1}}$, *is smaller than* $\left(\frac{\sigma-1}{\sigma}\right)^{\sigma-2}$. *If the profit share is smaller than* $1/e$, *consumption declines irrespective of the demand elasticity.*

Proof. The critical level of the share of exceptional outcomes in expected development profits before automation is $\left(1+\frac{\lambda}{\sigma\rho/\delta-1}\right)^{2-\sigma}$: consumption falls when the share is below this level (see (28)). Since $\rho = 1-\beta(1-\delta) \geq \delta$ and $0 < \lambda < 1$,

$$1 < \left(1+\frac{\lambda}{\sigma\rho/\delta-1}\right)^{\sigma-2} \leq \left(1+\frac{\lambda}{\sigma-1}\right)^{\sigma-2} < \left(\frac{\sigma}{\sigma-1}\right)^{\sigma-2} < e.$$

□

This corollary suggests that the introduction of AI automation may be welfare harming. Moreover, notice that the profit share condition is relatively loose. For a profit share of ordinary outcomes above $(e-1)/e \approx 63\%$, we obtain a consumption decline even when

the demand elasticity is extremely large. For a more common value of demand elasticities in models with a love of variety, for example, when $\sigma = 3$, it is sufficient that ordinary outcomes make up at least one third of expected profits of product development in order to make the AI automation welfare harming.[8]

4.3 Competitive Automation of Ordinary Success

This negative effect of AI automation worsens if AI products are subject to fiercer competition, as we will show next.

AI in production. We now assume that many AI suppliers compete for the production of a human idea. Therefore, AI products are sold without a markup (and without profits) at a price $1/z_A$, no higher than the low human supplier's marginal cost $1/z_L$. Displacement now only requires $z_A \geq z_L$ and, at equality, we let AI serve the market. The high-productivity human provider remains protected from AI and charges μ.

Labor demand from AI production. As in the previous subsection, average productivity is given by (20). However, what differs is the price level and the market share of AI products. Let ℓ_A^A again denote production labor per AI-supplied product. Relative CES demand and the two production technologies imply

$$\ell_A^A = \mu^\sigma z_A^{\sigma-1} \ell_A^H.$$

The term μ^σ captures that the automated goods and services are now sold in a competitive market and, thus, their suppliers charge no markup, while the suppliers of exceptional human products do. Compared to the common markup case, this distorts consumption and labor towards the AI-provided goods and services. The fraction of production labor employed by human-design suppliers, denoted by $1 - \lambda$, is therefore

$$1 - \lambda = \frac{s}{s + (1-s)\mu^\sigma z_A^{\sigma-1}} \tag{32}$$

and thus lower than when AI automated products also sell at a markup. As before, higher AI productivity lowers this share. (Note that with different markups, labor-share and market share differ).

Development of goods. The aggregate free-entry conditions remain as in the previous subsection, see (24) and (25) (left-hand side before substituting in prices). However, as is formalized in the following proposition, AI automation has a stronger impact on product development than in the previous subsection.

[8] Note that the previous example is obtained as the limit of $z_L \downarrow 0$. Then the share of expected profits from the exceptional outcomes tends to 1 and the savings in development costs always dominate.

PROPOSITION **5** (The loss of ordinary commercial returns)
Consider $\kappa_o = 0$. If AI is used to produce only ordinary goods, with productivity $z_A > z_L$, and charges no markup, then automation implies that

i) the mass of varieties has a more negative elasticity in AI productivity than in the case with equal markups. As a result, real wages and output fall when AI productivity increases.

ii) relative to the baseline case without AI, the real wage, output, the mass of varieties, and development investment are all lower, and

$$\frac{Y_A}{Y_b} = \left(1 - \frac{\lambda}{\sigma}\right) \frac{P_b}{P_A} < 1 \qquad \frac{M_A}{M_b} = \frac{I_A}{I_b} = (1-\lambda)\frac{P_b}{P_A} < 1, \tag{33}$$

$$\frac{P_b}{P_A} = \left(\frac{s}{\tilde{z}_b^{\sigma-1}}\right)^{\frac{1}{\sigma-2}} \left(1 - \frac{\lambda}{\sigma}\right)^{\frac{1}{\sigma-2}} < 1. \tag{34}$$

iii) consumption declines relative to baseline.

Proof. To express the implied mass of supplied goods, we can use the definition of λ, to write

$$s + (1-s)(\mu z_A)^{\sigma-1} = \frac{s}{1-\lambda}\left(1 - \frac{\lambda}{\sigma}\right).$$

With competitive AI supply, the price index becomes

$$P_A = \mu M_A^{\frac{1}{1-\sigma}} \left[s + (1-s)(\mu z_A)^{\sigma-1}\right]^{\frac{1}{1-\sigma}} = \mu \left[M_A \frac{s}{1-\lambda}\left(1 - \frac{\lambda}{\sigma}\right)\right]^{\frac{1}{1-\sigma}}.$$

Substituting the price index into the aggregate free-entry conditions yields the following equilibrium conditions before and after automation:

$$M_b = \left(\frac{L}{\sigma\rho\kappa_d}\right)^{\frac{\sigma-1}{\sigma-2}} \tilde{z}_b^{\frac{\sigma-1}{\sigma-2}}, \qquad M_A = \left(\frac{Ls\left(1 - \frac{\lambda}{\sigma}\right)}{\sigma\rho\kappa_d}\right)^{\frac{\sigma-1}{\sigma-2}} \left[\frac{s}{1-\lambda}\left(1 - \frac{\lambda}{\sigma}\right)\right]^{-1}. \tag{35}$$

As in the last subsection, after automation, only high quality products generate profits. Now, however, the real wage is not constant but rather decreases in AI productivity:

$$P_A^{-1} = \frac{1}{\mu}\left(\frac{Ls}{\sigma\rho\kappa_d}\right)^{\frac{1}{\sigma-2}} \left(1 - \frac{\lambda}{\sigma}\right)^{\frac{1}{\sigma-2}}$$

Parts (i) and (ii) follow from the preceding derivations. For part (iii), the consumption expressions before and after automation imply

$$\frac{C_A}{C_b} = \frac{P_b}{P_A}\left[1 - \frac{(\rho-\delta)\lambda}{\sigma\rho - \delta}\right] < 1.$$

The real-wage ratio is below one by part (ii), while the second factor is at most one because $\rho \geq \delta$: at a given real wage, the output loss weakly exceeds the savings in development goods. Consumption therefore falls, including in the patient limit $\rho = \delta$, when these two resource effects offset each other. □

Intuitively, cheaper AI products raise purchasing power at a given number of varieties, but they also draw demand away from human-developed products. The factor $\mu^{-1} < \left(1 - \frac{\lambda}{\sigma}\right) < 1$ reflects the fact that competitive AI producers capture a market share that is larger than what is warranted by their productivity alone. When it comes to product development, this overcompensates the positive partial effect $\frac{s}{1-\lambda}\left(1 - \frac{\lambda}{\sigma}\right)$ on the real wage. The lower expected returns to development lead to a decline in entry. This decline more than offsets the direct gain from lower prices. In equilibrium, the real wage falls as if market size had declined from L to $L(1 - \lambda/\sigma)$.

The real wage also provides a convenient expression for aggregate output. Holding the supplied products and their technologies fixed, the marginal contribution to composite output of a worker at a human producer is μ/P_A, whereas it is $1/P_A$ at an AI producer. Since aggregate production is homogeneous of degree one in labor inputs, output satisfies

$$Y_A = [\mu(1-\lambda) + \lambda]\frac{L}{P_A} = \mu\left(1 - \underbrace{\frac{\lambda}{\sigma}}_{\text{lower marginal product}}\right)\frac{L}{P_A}.$$

Here, the term $\frac{\lambda}{\sigma}$ reflects that the marginal product of AI goods is lower than the marginal product of human-design products due to the missing markup. Therefore, a marginal increase in AI's labor market share λ lowers output for a given price level by $\frac{\mu}{\sigma}\frac{L}{P_A}$.

Note that this output decrease actually exceeds the savings from a larger share of products being available without the need to pay more development costs as the success rate goes up. Free entry implies $\rho\kappa_d M_A = (\mu - 1)(1 - \lambda)L/P_A$. With replacement investment $I_A = \delta M_A$ and no operating costs, consumption is therefore

$$C_A = Y_A\left[1 - \frac{\delta}{\rho}\frac{1-\lambda}{\sigma - \lambda}\right] = \mu\frac{L}{P_A}\left[1 \quad \underbrace{-\frac{\delta}{\rho}\frac{1}{\sigma}}_{\substack{\text{baseline}\\ \text{development goods}}} \quad + \quad \underbrace{\frac{\delta}{\rho}\frac{\lambda}{\sigma}}_{\substack{\text{development}\\ \text{goods saved}}} \quad - \quad \underbrace{\frac{\lambda}{\sigma}}_{\substack{\text{lower}\\ \text{marginal product}}}\right].$$

As $\delta < \rho$, the fact that AI products have lower marginal products dominates. Increasing their market share, therefore, lowers consumption even at fixed real wages.

4.4 Pricing and equilibrium in the general model

We now generalize the competitive-automation economy of Section 4.3. Human productivity is drawn from G^H, and AI productivity from the conditional distribution $G^A(\cdot \mid z^H)$. Both draws are observed before production and remain fixed while the product is in fashion. We restore the common operating cost $\kappa_o > 0$, paid in goods by the supplier serving the market. Developers anticipate AI competition and receive no royalties when AI produces their designs. When a product falls out of fashion, neither technology can continue to supply it.

Pricing and production. As before, CES demand is $y = YP^{\sigma}p^{-\sigma}$. For pricing and development decisions, it is convenient to collect the aggregate terms in

$$X := YP^{\sigma-1}, \qquad y(\omega) = p(\omega)^{-\sigma}PX. \tag{36}$$

Holding X fixed keeps the profit schedule in goods units fixed at any given price and technology. This is what firms compare to the goods cost of development and operation. A supplier with productivity $z > 0$ earns a profit in goods of $Xp^{-\sigma}(p - 1/z) - \kappa_o$ if it serves the market. Its highest attainable profit is $X(z/\mu)^{\sigma-1}/\sigma - \kappa_o$. The minimum productivity that permits production is therefore

$$z^{\mathrm{be}}(X) := \mu \left(\frac{\sigma\kappa_o}{X}\right)^{1/(\sigma-1)}. \tag{37}$$

All potential AI suppliers of a product have access to the same AI draw. We assume contestable supply: a potential supplier can offer to serve the entire product market and incurs the operating cost only if it does so. Viable AI consequently offers the lowest price that covers variable and operating costs. Denoting this price by p_A, it satisfies

$$Xp_A^{-\sigma}\left(p_A - \frac{1}{z^A}\right) = \kappa_o, \qquad \frac{1}{z^A} < p_A \leq \frac{\mu}{z^A}. \tag{38}$$

The stated interval selects the lower break-even root. Unlike in Section 4.3, competitive AI generally charges more than marginal cost, because it must cover overhead. Its net profit remains zero. AI below the viability threshold imposes no price ceiling.

LEMMA **1** (Pricing and production)
A developed product is supplied if and only if $\max\{z^H, z^A\} \geq z^{\text{be}}(X)$*, with production at equality. If AI is not viable, a viable human charges* μ/z^H*. If AI is viable and* $z^A \geq z^H$*, AI supplies the product at* p_A*. If AI is viable and* $z^H > z^A$*, the human charges its unconstrained monopoly price when* $\mu/z^H \leq p_A$ *and otherwise limit-prices at* p_A*. That is, the human-design supplier sets*

$$p_H = \min\left\{\frac{\mu}{z^H}, p_A\right\}$$

and earns positive profits. At equal prices, a human earning strictly positive profits serves the market; otherwise viable AI serves it.

Proof. The maximum profit above gives the viability threshold. At the AI break-even price, the human's profit in goods is

$$Xp_A^{-\sigma}\left(\frac{1}{z^A} - \frac{1}{z^H}\right).$$

A more productive human can therefore cover the same overhead and earn a positive margin, whereas a less productive human cannot profitably match the AI price. The more productive human chooses its monopoly price unless the AI price ceiling binds. A product that neither technology can profitably supply at its monopoly price cannot be supplied at a lower price. □

Selection and average productivity. The productivity actually used for a supplied product is the higher of the two draws. Consequently, the supplied fraction of developed ideas and average supplier productivity are

$$\begin{aligned} s^* &:= \Pr\big(\max\{z^H, z^A\} \geq z^{\text{be}}(X)\big), \\ \tilde{z} &:= \big[\mathbb{E}\big[\max\{z^H, z^A\}^{\sigma-1} \mid \text{product supplied}\big]\big]^{1/(\sigma-1)}. \end{aligned} \tag{39}$$

As before, $M \equiv s^*K$ distinguishes supplied products from the stock of developed ideas. AI weakly improves the technology available for every idea. Nevertheless, average productivity among supplied products can fall if AI makes additional low-productivity ideas viable. Nor does a larger supplied fraction necessarily imply more products, because development responds to the returns that humans retain.

Development of goods. Let π^H denote the human developer's net profit in wage units, as in the baseline, and set it to zero when the human does not produce. When the

human supplies the product, $\pi^H/P = Xp_H^{-\sigma}(p_H - 1/z^H) - \kappa_o$. Free entry requires

$$\mathbb{E}\left[\frac{\pi^H}{P}\right] = \rho\kappa_d, \tag{40}$$

where expectations are over the human draw and the conditional AI draw. Assume $\mathbb{E}[\max\{z^H, z^A\}^{\sigma-1}] < \infty$ and $\Pr(z^H > z^A) > 0$. The latter condition allows some developers to retain positive returns; if AI always matches or exceeds human productivity, development has no private value, and there is no positive-entry equilibrium. Expected human profits in goods are continuous and strictly increasing in X wherever positive: an unconstrained human benefits from higher demand, while a human constrained by AI sells more at a positive cost advantage. These returns become arbitrarily large as X rises. Free entry, therefore, determines a unique positive X under the stated participation rules.

Output and equilibrium scale. Different suppliers now charge different markups, so average supplier productivity and the number of products alone no longer determine gross output. Individual labor demand is $\ell(\omega) = PXp(\omega)^{-\sigma}/z(\omega)$; relative labor inputs therefore follow from the product prices and technologies. Let $\bar{m}$ denote the average markup over marginal production cost, weighted by production labor:

$$\bar{m} := \frac{1}{L}\int_\Omega p(\omega)z(\omega)\ell(\omega)\,d\omega, \qquad PY \equiv \bar{m}L. \tag{41}$$

Every supplied product has a markup between one and μ, so $1 < \bar{m} \leq \mu$. With common markups, we recover $Y = \mu L/P$. With competitive AI, the average markup replaces μ in this relation. It measures the relation between gross output and the real wage, rather than allocation efficiency while holding the supplied products and their technologies fixed. Combining this relation with the definition of X gives

$$\frac{1}{P} = \left(\frac{\bar{m}L}{X}\right)^{1/(\sigma-2)}, \qquad Y = \frac{\bar{m}L}{P}. \tag{42}$$

Gross revenue pays production wages, operating costs, and developer profits. Since AI earns zero net profits and free entry gives total human profits of $\rho P\kappa_d K$, we have $(\bar{m}-1)L/P = K(\kappa_o s^* + \rho\kappa_d)$. Thus, the stock of developed ideas, the flow of development, and the mass of supplied products satisfy

$$K = \frac{L}{P}\frac{\bar{m}-1}{\kappa_o s^* + \rho\kappa_d}, \qquad I = \delta K, \qquad M \equiv s^* K. \tag{43}$$

These expressions provide the aggregate counterpart of the entry conditions in the examples. Free entry determines X; pricing and selection then determine s^* and $\bar{m}$; the real wage and the stock of ideas follow from the aggregate equations.

Consumption and the use of goods. The resource constraint continues to distinguish gross output from consumption. Substituting equilibrium development into $C = Y - K(\kappa_o s^* + \delta\kappa_d)$ gives

$$\begin{aligned} C &= \frac{L}{P}\left[\bar{m} - (\bar{m}-1)\frac{\kappa_o s^* + \delta\kappa_d}{\kappa_o s^* + \rho\kappa_d}\right] \\ &= \frac{L}{P}\left[1 + \frac{(\rho-\delta)\kappa_d(\bar{m}-1)}{\kappa_o s^* + \rho\kappa_d}\right]. \end{aligned} \tag{44}$$

This is the same resource comparison as in Section 4.3. Holding the real wage and the supplied fraction fixed, a lower average markup reduces gross output but also reduces the goods absorbed by development and operation. The latter saving equals the fraction $(\kappa_o s^* + \delta\kappa_d)/(\kappa_o s^* + \rho\kappa_d)$ of the output reduction. Since $\rho \geq \delta$, the saving cannot exceed that reduction. In the patient limit, the two effects cancel at a given real wage and $C = L/P$; the equilibrium response of the real wage remains decisive.

PROPOSITION **6** (Introducing automation)
Under the pricing and participation assumptions above, with $\sigma > 2$, finite productivity moments, and $Pr(z^H > z^A) > 0$ and relative to the economy without AI, competitive automation weakly increases the supplied fraction, $s_A^* \geq s_b^*$, but weakly reduces the stock of developed ideas, development investment, the real wage, gross output, and stationary consumption. Specifically,

$$\frac{P_b}{P_A} = \left[\frac{\bar{m}_A}{\mu}\frac{X_b}{X_A}\right]^{1/(\sigma-2)} \leq 1, \qquad \frac{Y_A}{Y_b} = \frac{\bar{m}_A}{\mu}\frac{P_b}{P_A} \leq 1, \tag{45}$$

$$\frac{K_A}{K_b} = \frac{I_A}{I_b} = \frac{P_b}{P_A}\frac{\bar{m}_A - 1}{\mu - 1}\frac{\kappa_o s_b^* + \rho\kappa_d}{\kappa_o s_A^* + \rho\kappa_d} \leq 1, \qquad \frac{M_A}{M_b} = \frac{s_A^*}{s_b^*}\frac{K_A}{K_b}, \tag{46}$$

and

$$\frac{C_A}{C_b} = \frac{P_b}{P_A}\frac{1 + \dfrac{(\rho-\delta)\kappa_d(\bar{m}_A - 1)}{\kappa_o s_A^* + \rho\kappa_d}}{1 + \dfrac{(\rho-\delta)\kappa_d(\mu - 1)}{\kappa_o s_b^* + \rho\kappa_d}} \leq 1. \tag{47}$$

The mass of supplied products can rise or fall. The declines in development, the real wage, output, and consumption are strict if automation raises X or lowers $\bar{m}$ below μ.

Proof. At a given X, AI competition cannot increase a human developer's profit. Free entry, therefore, requires $X_A \geq X_b$. A higher X lowers the productivity needed to cover overhead, and AI expands the technologies available for each idea, implying $s_A^* \geq s_b^*$. Without AI, every supplier charges μ; with competitive AI, no supplier charges more, so $\bar{m}_A \leq \mu$. Substituting these comparisons into the aggregate expressions proves the results, using $\rho \geq \delta$ for consumption. A larger supplied fraction and a smaller stock of developed ideas have opposing effects on M. □

In the two-point example of Section 4.3, $\kappa_o = 0$, $X_b/X_A = s/[s + (1-s)z_L^{\sigma-1}]$, and $\bar{m}_A = \mu(1 - \lambda/\sigma)$. The consumption comparison, therefore, reduces to

$$\frac{C_A}{C_b} = \frac{P_b}{P_A}\left[1 - \frac{(\rho - \delta)\lambda}{\sigma\rho - \delta}\right].$$

The general result preserves the same logic. AI can improve production technologies while lowering the returns that support development. The resulting reduction in development saves resources, but it does not overturn the consumption loss relative to no AI under competitive supply. The comparison concerns stationary equilibria and does not include consumption along the transition.

Further improvements in AI. The proposition compares competitive automation with no AI. It does not imply that every improvement in AI reduces consumption in an economy already using it. For two such economies, indexed by 1 and 2, the patient-limit comparison is

$$\frac{C_2}{C_1} = \left[\frac{\bar{m}_2}{\bar{m}_1}\frac{X_1}{X_2}\right]^{1/(\sigma-2)}. \tag{48}$$

An improvement that weakly raises AI productivity at every conditional percentile cannot lower the demand required to sustain development, so $X_2 \geq X_1$. Its effect on the average markup is nevertheless ambiguous. For example, AI can make previously unserved ideas with zero human productivity just viable. These additional suppliers charge μ, with their entire margin covering overhead, and leave human profits at a given X unchanged. If existing AI suppliers charge lower markups, the new products raise $\bar{m}$ and consumption rises in the patient limit. This composition effect cannot raise the average markup above the no-AI benchmark, which is why the comparison in Proposition 6 remains valid.

5 Extensions and Robustness

We next discuss what happens when augmentation and automation occur jointly, how results change when we assume development costs are paid in labor, and what a dynamic version with a transition path would look like. Details for all three extensions can be found in Appendices D - F.

5.1 Augmentation and Automation jointly

The previous two sections described the effects of augmentation and automation separately. In this section, we analyze the combination of the opposing effects.

5.1.1 Results for the general model

When the productivity of human supplied products changes alongside AI products gaining market shares that do not reward developers, the two aspects of AI introduction interact. Appendix D formalizes the two roles of AI in the general model. The joint consumption effect can be written as the gain from augmentation multiplied by the effect of automation in the augmented economy. However, the second factor generally changes with augmentation because higher human productivity affects which developers retain their markets, whether AI constrains their prices, and the allocation of production. Having said this, with the conditional distribution of AI productivity held fixed, sufficiently strong uniform augmentation of human production eventually raises consumption above its no-AI level.

5.1.2 An Illustrative Quantification

We now return to our two-outcome (ordinary and exceptional) special case to illustrate the key points stemming from the combined effect of AI on augmentation and automation. The two-outcome example is particularly convenient because we obtain a unique augmentation threshold and explicit bounds. Specifically, restoring development profits to the level with ordinary outcomes contributing is necessary, but stronger augmentation is needed to overcome the effects of unequal markups. Thus, the comparison with no AI depends on the strength of augmentation, even though competitive automation continues to reduce consumption relative to the augmented economy without automation.

AI in production. To illustrate how much augmentation is needed to offset automation's effect on the real wage, we return to the two viable outcomes of Sections 4.2 and 4.3, with zero operating costs. We normalize $z_H = 1$ and express the thought experiment in terms of the labor share of "ordinary outcomes". This avoids specifying the ordinary-product productivity z_L and probability s separately. It is a sufficient statistic in the case where AI products are sold with a markup too.

Uniform augmentation first multiplies both human productivities by A. AI then supplies ordinary outcomes, while exceptional outcomes remain protected, and the developers of automated products receive no royalties. For this exercise, AI productivity scales with A as well, leaving relative human and AI productivities unchanged. This is the special case in which the augmentation gain and the subsequent automation effect separate exactly. We compare stationary equilibria, with developers anticipating both changes in the AI economy.

Augmentation needed to restore the baseline real wage. Figure 4 plots the productivity increase, in percent, required to restore the original real wage. The horizontal

axis is the pre-AI labor share of ordinary products, $q := 1 - s/[s + (1-s)z_L^{\sigma-1}]$, which is also their contribution to expected development profits as a share of the total. Panel (a) retains the common markup μ, as in Section 4.2. Panels (b) and (c) instead have competitive AI supply, with $z_A = A\mu z_L$ and $z_A = Az_H$, respectively. Writing λ for AI's labor share after automation, the required augmentation satisfies

$$A = \begin{cases} (1-q)^{-1/(\sigma-1)}, & \text{equal markups,} \\ \left[(1-q)(1-\lambda/\sigma)\right]^{-1/(\sigma-1)}, & \text{competitive AI.} \end{cases} \tag{49}$$

The first expression replaces the ordinary returns lost to development. The second requires additional augmentation because unequal markups draw labor toward AI products and reduce the real wage further. The relevant AI employment share, λ, is endogenous and generally exceeds the pre-AI share, q. For case (b), λ is a function of demand elasticities and the pre-automation labor share q: $\frac{\lambda}{1-\lambda} = \mu^{2\sigma-1}\frac{q}{1-q}$. For case (c), we need to specify z_L and set this to 1/2 (though numerically the results change little when z_L changes).

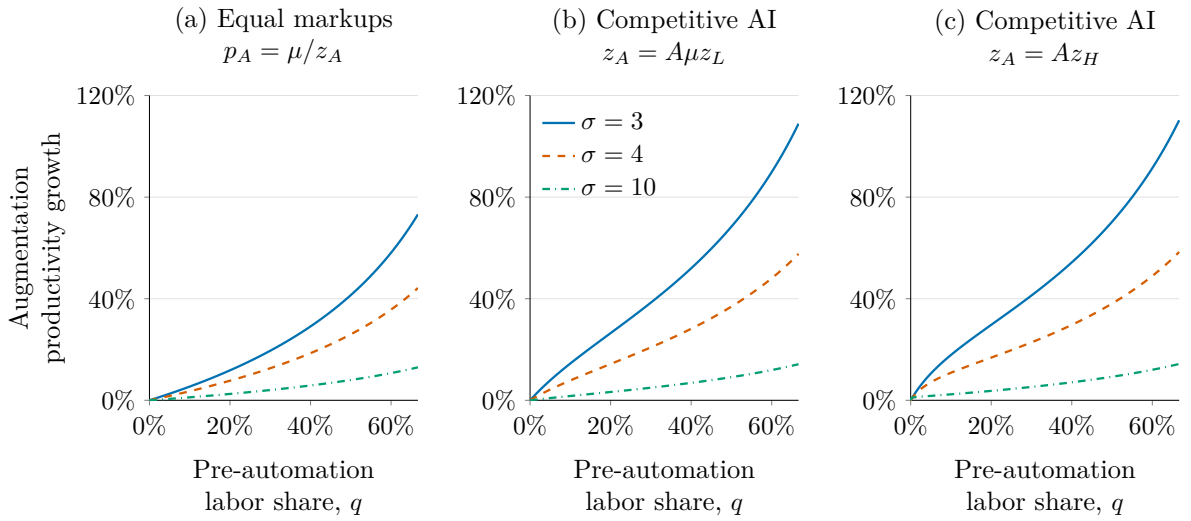
Illustrative results. We focus on $\sigma = 4$, a conservative choice for the required augmentation relative to the baseline calibrations. [9] A higher elasticity makes varieties closer substitutes and lowers the augmentation needed at a given pre-AI labor share. The curves for $\sigma = 3$ and $\sigma = 10$ show the sensitivity to this parameter.

Consider ordinary products accounting for one third of employment and expected development profits before AI, $q = 1/3$. At $\sigma = 4$ and $\frac{z_L}{z_H} = \frac{1}{2}$, this corresponds to an ordinary-outcome probability $(1-s)$ of 80%. With equal markups, augmentation of 14.5% restores the real wage: the increase in exceptional productivity raises its contribution to expected profits enough to replace the lost ordinary returns. With competitive AI, the required productivity increase through augmentation rises to 23.2% when $z_A = A\mu z_L$ and 25.0% when $z_A = Az_H$. If augmentation of the remaining human-design products falls short of these thresholds, the real wage declines.

These numbers also restore consumption under competitive AI in the patient limit, when $\rho = \delta$ and $C = L/P$. For $\rho > \delta$, stronger augmentation is needed to restore consumption. With equal markups, by contrast, restoring the real wage already leaves consumption above its original level because fewer goods are used for development. The real-wage and consumption thresholds, therefore, depend differently on how goods are allocated between consumption and development.

[9] For example, Bachmann et al. (2026) use an elasticity of 3.5 for the private non-primary sector and Bilbiie et al. (2012) use 3.8.

Figure 4: Augmentation required to preserve the real wage



Notes: The vertical axis shows the required augmentation productivity gain in percent to keep the real wage unchanged relative to the baseline economy. The horizontal axis is the pre-AI labor share of ordinary products, $q := 1 - s/[s + (1-s)z_L^{\sigma-1}]$. We set $z_H = 1$ (and specify $z_L = 1/2$ for (c)), and $\kappa_o = 0$, and vary q for each elasticity. Uniform augmentation multiplies both human productivities by A, after which AI supplies only ordinary outcomes. Panel (a) has equal markups and $z_A > A\mu z_L$. Panels (b) and (c) have competitive AI, with $z_A = A\mu z_L$ and $z_A = Az_H$, respectively. Developers receive no royalties from AI production; in panel (a), AI profits accrue to households as dividends. All comparisons are with the original stationary equilibrium before augmentation and automation.

5.2 Development costs in terms of labor

Appendix E considers development costs paid in labor. With zero operating costs, equilibrium development and the allocation of labor between development and production are independent of the productivity distribution. Uniform augmentation, therefore, raises consumption proportionally to productivity, without the additional gain through greater product variety obtained when development costs are paid in goods. Competitive automation of ordinary outcomes, however, still reduces development, product variety, and stationary consumption. Fewer development attempts release workers into production, but this reallocation does not compensate for the loss of ordinary development returns and the effects of unequal markups.

5.3 Transitions

Appendix F returns to development costs in goods and derives the transition following an unexpected introduction of competitive AI in the two-viable-outcome model. The simple model has an arguably extreme transition: after a surprising once-and-for-all introduction of AI automation, development stops while the surviving stock of designs exceeds its new stationary level. Initially, AI can therefore draw on previously developed products, potentially even making their production more efficient, while reduced development makes more goods available for consumption. These gains can outweigh the subsequent losses

from fewer varieties: at any fixed discount factor, sufficiently productive AI or sufficiently long-lived product ideas raise discounted welfare despite lowering stationary consumption. Conversely, for any fixed, finite AI productivity, sufficiently patient households experience a welfare loss.

6 Discussion and Conclusion

The present paper discusses the effects of AI in a framework of creative destruction, where the development of products and ideas is pivotal to an economy's productivity. If the introduction of AI augments the productivity of existing production processes, then we find that the aggregate equilibrium productivity gains can be expected to be even stronger than the productivity gains at the microeconomic level if all firms benefit equally. If augmentation disproportionately raises the productivity of initially more productive firms, then the gains are even higher.

However, if AI introduces products that substitute previously human-provided products and services, then this can have drastically negative long-run welfare effects. As human developers expect to receive lower compensation for their development efforts, they will exert less effort to develop ideas for new products. As a result, fewer ideas are developed, affecting the varieties available to consumers. The economy becomes less efficient, causing consumption to fall in the long run.

Strictly speaking, our model is a model of firms developing products. Yet, we expect the basic mechanisms to carry over to other environments. One example could be education, when its returns are uncertain. Ordinary jobs contribute to the expected payoff from training, while exceptional realizations can generate major discoveries or creative achievements. If AI erodes the ordinary returns, fewer people may enter these careers, reducing the number of exceptional outcomes. Consequently, AI may give rise to "missing generations" of new and revolutionary ideas of star architects, engineers, or scientists. The danger is that this will impoverish consumers in the long run.

References

Acemoglu, Daron, Dingwen Kong, and Asuman Ozdaglar (2026) "AI, Human Cognition and Knowledge Collapse," Working Paper34910, National Bureau of Economic Research, 10.3386/w34910.

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

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li (2023) "A Theory of Falling Growth and Rising Rents," *Review of Economic Studies*, 90 (6), 2675–2702, 10.1093/restud/rdad016.

Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones (2019) "Artificial Intelligence and Economic Growth," in Ajay Agrawal, Joshua Gans, and Avi Goldfarb eds. *The Economics of Artificial Intelligence: An Agenda*, 237–282, Chicago: University of Chicago Press, https://www.nber.org/books-and-chapters/economics-artificial-intelligence-agenda/artificial-intelligence-and-economic-growth.

Aguiar, Luis and Joel Waldfogel (2018) "Quality Predictability and the Welfare Benefits from New Products: Evidence from the Digitization of Recorded Music," *Journal of Political Economy*, 126 (2), 492–524, 10.1086/696229.

American Federation of Musicians (n.d.) "125 Years: Musicians Staying Stronger Together," American Federation of Musicians, https://www.afm.org/125-years-afm-history/, Accessed 31 October 2026.

Arrow, Kenneth J. (1962) "Economic Welfare and the Allocation of Resources for Invention," in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, 609–626, Princeton, NJ: Princeton University Press, https://www.nber.org/system/files/chapters/c2144/c2144.pdf.

Asriyan, Vladimir, Hugo Reichardt, and Sandro Shelegia (2026) "Immiserizing Automation," CEPR Discussion Paper21872, Centre for Economic Policy Research, https://cepr.org/publications/dp21872.

Atkeson, Andrew and Ariel Burstein (2010) "Innovation, Firm Dynamics, and International Trade," *Journal of Political Economy*, 118 (3), 433–484, 10.1086/653690.

Bachmann, Rüdiger, Christian Bayer, Heiko Stüber, and Felix Wellschmied (2026) "Monopsony Makes Firms Not Only Small but Also Unproductive: Why East Germany

Has Not Converged," *The Review of Economic Studies*, 10.1093/restud/rdag033, Advance online publication.

Baldwin, Richard E. and Frédéric Robert-Nicoud (2008) "Trade and Growth with Heterogeneous Firms," *Journal of International Economics*, 74 (1), 21–34, 10.1016/j.jinteco.2007.05.004.

Baumol, William J. and Hilda Baumol (1994) "On the Economics of Musical Composition in Mozart's Vienna," *Journal of Cultural Economics*, 18, 171–198, 10.1007/BF01080225.

Bazzichi, Emanuele, Massimo Riccaboni, and Fulvio Castellacci (2026) "Bridging Distant Ideas: The Impact of AI on R&D and Recombinant Innovation," Working Paper2604.02189, arXiv, 10.48550/arXiv.2604.02189.

Bilbiie, Florin O., Fabio Ghironi, and Marc J. Melitz (2012) "Endogenous Entry, Product Variety, and Business Cycles," *Journal of Political Economy*, 120 (2), 304–345, 10.1086/665825.

Bonfiglioli, Alessandra, Rosario Crinò, and Gino Gancia (2018) "Betting on Exports: Trade and Endogenous Heterogeneity," *Economic Journal*, 128 (609), 612–651, 10.1111/ecoj.12408.

Dhingra, Swati and John Morrow (2019) "Monopolistic Competition and Optimum Product Diversity under Firm Heterogeneity," *Journal of Political Economy*, 127 (1), 196–232, 10.1086/700732.

Di Licosa, Martina and Matt Craig (2025) "The Highest-Paid Musicians of 2025," Forbes, December, https://www.forbes.com/sites/martinadilicosa/2025/12/30/the-highest-paid-musicians-of-2025/, Published 30 December 2025; updated 19 May 2026. Accessed 31 October 2026.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu (2023) "How Costly Are Markups?" *Journal of Political Economy*, 131 (7), 1619–1675, 10.1086/722986.

Gallini, Nancy T. (1992) "Patent Policy and Costly Imitation," *The RAND Journal of Economics*, 23 (1), 52–63, 10.2307/2555432.

Garicano, Luis and Luis Rayo (2025) "Training in the Age of AI: A Theory of Career Viability," Discussion Paper20634, Centre for Economic Policy Research, https://cepr.org/publications/dp20634, Revised March 2026.

Gilbert, Richard and Carl Shapiro (1990) "Optimal Patent Length and Breadth," *The RAND Journal of Economics*, 21 (1), 106–112, 10.2307/2555497.

Ide, Enrique (2025) "Automation, AI, and the Intergenerational Transmission of Knowledge," Working Paper2507.16078, arXiv, 10.48550/arXiv.2507.16078, Revised June 2026.

Intel Corporation (n.d.) "Announcing a New Era of Integrated Electronics: The Intel 4004," Intel History, https://www.intel.com/content/www/us/en/history/virtual-vault/articles/the-intel-4004.html, Accessed 31 October 2026.

Katsamakas, Evangelos and J. Manuel Sanchez-Cartas (2024) "Generative Artificial Intelligence, Content Creation, and Platforms," *Journal of Industry, Competition and Trade*, 24 (1), 10.1007/s10842-024-00430-9, Article 19.

Klemperer, Paul (1990) "How Broad Should the Scope of Patent Protection Be?" *The RAND Journal of Economics*, 21 (1), 113–130, 10.2307/2555498.

Korinek, Anton, Charles I. Jones, Szymon Sacher, Tess Cotter, and Peter McCrory (2026) "Economic Scenarios for Transformative AI," Working Paper2026-02, The Anthropic Institute, https://www-cdn.anthropic.com/files/4zrzovbb/website/cf58f84d46a4a76bf5a5b039ac695fba6b80041c.pdf.

Melitz, Marc J. (2003) "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71 (6), 1695–1725.

Nordhaus, William D. (1969) *Invention, Growth, and Welfare: A Theoretical Treatment of Technological Change*, Cambridge, MA: The MIT Press, https://mitpress.mit.edu/9780262140072/invention-growth-and-welfare/.

Rosen, Sherwin (1981) "The Economics of Superstars," *American Economic Review*, 71 (5), 845–858, https://www.jstor.org/stable/1803469.

Social Security Administration (n.d.) "National Average Wage Index," Office of the Chief Actuary, online data series, https://www.ssa.gov/OACT/COLA/AWI.html, Includes 2025 average annual wage data. Accessed 31 October 2026.

Šupka, Ondřej (n.d.) "Antonín Dvořák Biography," Antonin-dvorak.cz, Akademie klasické hudby, https://www.antonin-dvorak.cz/en/life/biography/, English translation by Karolina Hughes. Accessed 31 October 2026.

Waldoff, Jessica (2018) "Mozart and Finances," in Simon P. Keefe ed. *Mozart in Context*, Chap. 20, 170–180: Cambridge University Press, 10.1017/9781316848487.021.

Appendix

A Selection and Aggregate Scale **36**

B Proof of Proposition 2 **37**

C Existence and Uniqueness under Competitive Automation **39**

D Augmentation and Automation Together **39**

E Product Development with Labor **44**

F Transition after Competitive Automation **49**

A Selection and Aggregate Scale

This appendix collects the equilibrium expressions used in the proof of nonuniform augmentation. Development and operating costs are paid in goods, so all labor is used in production. The cutoff derivation assumes $\kappa_o > 0$; the zero-overhead examples are derived in the main text.

The operating cutoff. From monopoly pricing and the zero-profit condition (3), a producing firm's profit in goods is

$$\frac{\pi(z)}{P} = \kappa_o \left[\left(\frac{z}{z^*} \right)^{\sigma-1} - 1 \right].$$

Consequently, free entry (6) determines the cutoff through

$$\kappa_o \int_{z^*}^{\infty} \left[\left(\frac{z}{z^*} \right)^{\sigma-1} - 1 \right] dG(z) = \rho \kappa_d. \tag{50}$$

The left-hand side is continuous and strictly decreasing in z^* wherever positive. For a continuous distribution, its derivative is

$$-\frac{(\sigma-1)\kappa_o}{z^*} \int_{z^*}^{\infty} \left(\frac{z}{z^*} \right)^{\sigma-1} dG(z) < 0.$$

With a finite, positive productivity moment, expected profits tend to infinity as the cutoff approaches zero and to zero as it approaches infinity. The cutoff is therefore unique. The same argument applies to the augmented distribution G_α.

The real wage and development. Substituting $Y = \mu L/P$ from (8) into the zero-profit condition yields

$$\kappa_o = \frac{L(z^*)^{\sigma-1} P^{\sigma-2}}{\sigma \mu^{\sigma-2}}.$$

Thus, the real wage can be expressed through the productivity of the marginal firm. Evaluated in the augmented economy, this gives

$$\frac{1}{P_\alpha} = \left[\frac{L(z_\alpha^*)^{\sigma-1}}{\sigma \kappa_o \mu^{\sigma-2}} \right]^{1/(\sigma-2)}. \tag{51}$$

To obtain the stock of developed ideas, multiply (9) by $M_\alpha = s_\alpha^* K_\alpha$ and use $Y_\alpha = \mu L/P_\alpha$. This gives

$$K_\alpha = \frac{L}{(\sigma-1)P_\alpha(\kappa_o s_\alpha^* + \rho\kappa_d)}, \qquad M_\alpha \equiv s_\alpha^* K_\alpha, \qquad I_\alpha = \delta K_\alpha. \tag{52}$$

Finally, substituting these expressions into the goods resource constraint gives

$$C_\alpha = \frac{L}{P_\alpha}\left[1 + \frac{(\rho-\delta)\kappa_d}{(\sigma-1)(\kappa_o s^*_\alpha + \rho\kappa_d)}\right]. \tag{53}$$

The corresponding baseline expressions follow by replacing the subscript α with b. These equations distinguish the stock of developed ideas, K, from the mass of supplied products, M, and gross output, Y, from consumption, C.

B Proof of Proposition 2

The cutoff statements concern positive operating costs. For the strict comparisons with uniform augmentation, we take G to be a continuous, nondegenerate distribution of positive productivities. Without sufficient variation in productivity and augmentation, the corresponding comparisons are weak. All relevant productivity moments are finite, as assumed in the main text.

Selection and the effective cutoff. Let u^*_α denote the underlying productivity of the marginal producer, so $z^*_\alpha = \alpha(u^*_\alpha)u^*_\alpha$. Because augmentation preserves productivity ordering, (50) can be written as

$$\kappa_o \int_{u^*_\alpha}^{\infty} \left[\left(\frac{\alpha(u)u}{\alpha(u^*_\alpha)u^*_\alpha}\right)^{\sigma-1} - 1\right] dG(u) = \rho\kappa_d.$$

At the baseline underlying cutoff z^*_b, increasing augmentation, $\alpha'(u) > 0$, raises productivity ratios above the margin and therefore raises expected profits. The underlying cutoff must rise to restore free entry, implying $s^*_\alpha \leq s^*_b$. Decreasing augmentation, $\alpha'(u) < 0$, reverses both comparisons. The supplied-fraction inequalities are strict when the distribution assigns positive probability between the underlying cutoffs.

The effective productivity of the marginal firm rises in both cases. To see this, evaluate expected profits under G_α at the baseline effective cutoff z^*_b. Since $\alpha(u) > 1$, every initially profitable producer has higher productivity relative to this cutoff and earns higher profits in goods. Expected profits exceed $\rho\kappa_d$, so the effective cutoff must increase. Equation (51) then implies a higher real wage.

Development, supplied products, and consumption. When $\alpha'(u) > 0$, the higher real wage and the weakly smaller supplied fraction both raise K_α in (52). Development investment therefore rises, while the effects on $M_\alpha = s^*_\alpha K_\alpha$ can oppose each other. The smaller supplied fraction also reinforces the real-wage gain in (53), so consumption rises.

When $\alpha'(u) < 0$, the higher real wage and the weakly higher ratio $s^*_\alpha/(\kappa_o s^*_\alpha + \rho\kappa_d)$

imply more supplied products. Their effects on the stock of developed ideas can oppose each other. In the consumption expression, the higher supplied fraction works against the real-wage gain, but this opposing term vanishes as $\beta \to 1$ and $\rho \to \delta$. Continuity therefore gives a consumption gain for discount factors sufficiently close to one.

Comparison at the same mean productivity. Compare effective productivity $\alpha(u)u$ with uniform augmentation Au, where A is defined in (15). The two transformed distributions have the same mean. If $\alpha'(u) > 0$, the difference $[\alpha(u) - A]u$ is negative for low draws and positive for high draws. Since both transformations preserve the ordering of draws, G_α is a mean-preserving spread of the uniformly augmented distribution. If $\alpha'(u) < 0$, the signs reverse and G_α is a mean-preserving compression.

At a candidate effective-productivity cutoff $t > 0$, expected profits in goods are

$$\kappa_o \int_0^\infty \max\left\{\left(\frac{z}{t}\right)^{\sigma-1} - 1, 0\right\} dG_\alpha(z).$$

The profit function is convex in productivity and strictly convex above the cutoff because $\sigma > 2$. A spread therefore raises expected profits relative to uniform augmentation, while a compression lowers them. Under the stated nondegeneracy and strict monotonicity conditions, the comparisons are strict: augmentation changes productivity on a positive-probability set of profitable upper-tail draws. Uniform augmentation satisfies free entry at $t = Az_b^*$, so the monotonicity of expected profits in the cutoff gives

$$z_\alpha^* > Az_b^* \quad \text{if } \alpha'(u) > 0, \qquad z_\alpha^* < Az_b^* \quad \text{if } \alpha'(u) < 0.$$

Combining (51) and (53) gives the exact consumption comparison

$$\frac{C_\alpha/C_b}{A^{(\sigma-1)/(\sigma-2)}} = \left(\frac{z_\alpha^*}{Az_b^*}\right)^{(\sigma-1)/(\sigma-2)} \frac{1 + \dfrac{(\rho-\delta)\kappa_d}{(\sigma-1)(\kappa_o s_\alpha^* + \rho\kappa_d)}}{1 + \dfrac{(\rho-\delta)\kappa_d}{(\sigma-1)(\kappa_o s_b^* + \rho\kappa_d)}}. \tag{54}$$

In case (a), $\alpha'(u) > 0$, the cutoff factor exceeds one and the remaining fraction is at least one because $s_\alpha^* \le s_b^*$ and $\rho \ge \delta$. In case (b), $\alpha'(u) < 0$, the cutoff factor is below one and the remaining fraction is at most one. This proves the comparisons with uniform augmentation for every discount factor, including the patient limit. □

C Existence and Uniqueness under Competitive Automation

This appendix supplies the monotonicity argument underlying free entry in Section 4.4. Here $X = YP^{\sigma-1}$ is demand relative to the goods cost of development and operation. We keep the pricing and participation rules of Lemma 1 and the assumptions $\mathbb{E}[\max\{z^H, z^A\}^{\sigma-1}] < \infty$ and $\Pr(z^H > z^A) > 0$.

For a fixed AI productivity, the lower break-even price p_A in (38) weakly decreases as X increases. Whenever a more productive human is constrained by this price, its profit in goods is

$$\frac{\pi^H}{P} = X p_A^{-\sigma} \left(\frac{1}{z^A} - \frac{1}{z^H} \right) > 0.$$

Both higher demand and a lower AI price raise the quantity sold at the human's positive cost advantage. Hence, profits increase strictly with X. An unconstrained human earns $X(z^H/\mu)^{\sigma-1}/\sigma - \kappa_o$, which also increases strictly. A human with $z^H \leq z^A$ never earns positive profits: whenever its technology could operate, AI is already viable. Profits are continuous when the human starts producing or the AI price ceiling begins to bind. In particular, when AI first becomes viable, $p_A = \mu/z^A$ exceeds the monopoly price of any strictly more productive human.

These observations imply that expected human profits in goods are continuous and strictly increasing wherever positive. On a bounded interval of X, profits are bounded above by $X(z^H/\mu)^{\sigma-1}/\sigma$, whose expectation is finite. Expected profits therefore tend to zero as $X \downarrow 0$. For every draw with $z^H > z^A$, profits grow without bound as $X \to \infty$: either the human remains unconstrained or p_A approaches $1/z^A$ while its positive cost advantage remains. Since this event has positive probability, expected profits also become unbounded. Free entry (40) consequently determines a unique positive X. The aggregate equations in Section 4.4 then determine the remaining equilibrium quantities.

D Augmentation and Automation Together

We now allow AI to augment human production and to compete with human developers at the same time. The distinction is between the productivity improvement available to a developer and the share of the resulting returns that the developer retains. Higher human productivity can restore incentives to develop products, but its aggregate effect depends on how it changes competition with AI.

D.1 The general model

Let u be the underlying human productivity draw from G. Augmentation gives the human productivity $z^H = \alpha(u)u$, with $\alpha(u) \geq 1$, while AI draws z^A from $G^A(\cdot \mid u)$. Both draws are observed before production. Conditioning AI productivity on the underlying human draw makes clear what is held fixed when augmentation changes. Unless stated otherwise, stronger augmentation leaves this conditional AI distribution unchanged. Development and operating costs remain in goods, and competitive AI earns zero net profits and pays no royalties. We impose the moment and participation conditions of Section 4.4 on the resulting pair of productivities.

Subscript b continues to denote the baseline without AI, subscript α denotes augmentation without automation, and subscript A now denotes the economy with both. The pricing rule and aggregate equations in Section 4.4 apply after replacing human productivity by $\alpha(u)u$.

PROPOSITION **7** (Augmentation with competitive automation)
Under the maintained assumptions:

(i) Relative to the augmented economy without automation, competitive automation weakly lowers stationary consumption: $C_A \leq C_\alpha$. The comparison with the original baseline satisfies

$$\frac{C_A}{C_b} = \left[\frac{\bar{m}_A}{\mu}\frac{X_b}{X_A}\right]^{1/(\sigma-2)} \frac{1 + \dfrac{(\rho-\delta)\kappa_d(\bar{m}_A - 1)}{\kappa_o s_A^* + \rho\kappa_d}}{1 + \dfrac{(\rho-\delta)\kappa_d(\mu - 1)}{\kappa_o s_b^* + \rho\kappa_d}}. \tag{55}$$

(ii) A consumption gain requires $X_A < X_b$. In the patient limit, the necessary and sufficient condition is

$$\frac{X_A}{X_b} < \frac{\bar{m}_A}{\mu}. \tag{56}$$

For any discount factor, a sufficient condition is

$$\frac{X_A}{X_b} < \frac{1}{\mu}\left[\frac{(\sigma-1)\rho}{\sigma\rho - \delta}\right]^{\sigma-2}. \tag{57}$$

(iii) If augmentation raises every positive human productivity by the same factor a, holding the conditional distribution of AI productivity fixed, sufficiently large a raises stationary consumption above C_b.

Proof. For part (i), use the augmented human distribution G_α as the no-automation benchmark in Proposition 6. At a given X, AI cannot increase the developer's retained profits. Consequently, $X_A \geq X_\alpha$, $s_A^* \geq s_\alpha^*$, and $\bar{m}_A \leq \mu$, implying $C_A \leq C_\alpha$. Comparing the general real-wage and consumption expressions with the original baseline gives (55).

For part (ii), suppose $X_A \geq X_b$. Since $\alpha(u)u \geq u$ and AI provides an additional technology, the lower viability threshold then implies $s_A^* \geq s_b^*$. Together with $\bar{m}_A \leq \mu$, both factors in (55) are at most one. Consumption therefore cannot rise. When $\rho = \delta$, the second factor equals one, giving (56).

For the sufficient condition, the general consumption equation implies $C_A \geq L/P_A$, while

$$C_b \leq \frac{L}{P_b}\left[1 + \frac{\rho - \delta}{\rho(\sigma - 1)}\right] = \frac{L}{P_b}\frac{\sigma\rho - \delta}{(\sigma - 1)\rho}.$$

Since $\bar{m}_A \geq 1$, the real-wage comparison gives

$$\frac{C_A}{C_b} \geq \left[\frac{X_b}{\mu X_A}\right]^{1/(\sigma-2)} \frac{(\sigma - 1)\rho}{\sigma\rho - \delta}.$$

Condition (57) makes this lower bound exceed one.

For part (iii), fix any $X > 0$. As $a \to \infty$, the human monopoly price $\mu/(au)$ approaches zero for each $u > 0$, eventually falling below any viable AI price. Its profit in goods then equals $X(au/\mu)^{\sigma-1}/\sigma - \kappa_o$ and becomes unbounded. The baseline has positive productive draws, so expected retained profits at this fixed X also become unbounded. Free entry thus requires $X_A(a) \to 0$. Eventually (57) holds. This proves eventual recovery of consumption; it does not require consumption to be monotone in augmentation throughout the general model. □

Part (ii) gives an incentive condition that can be checked before solving for aggregate scale. If, evaluated at baseline demand relative to goods costs, X_b, expected retained developer profits under both technologies do not exceed $\rho\kappa_d$, then their joint introduction cannot raise consumption. For a gain, the improvement in human productivity must do more than compensate developers for the profits lost to AI at this fixed demand level. In the patient limit, it must lower the demand needed to sustain development enough to offset the decline in the average markup. Positive operating costs also make changes in the supplied fraction matter away from that limit.

Sequential comparisons and interaction. The exact factorization

$$\frac{C_A}{C_b} = \frac{C_\alpha}{C_b}\frac{C_A}{C_\alpha}$$

allows us to organize the analysis as augmentation followed by automation. However, the second factor is the effect of automation in the augmented economy. Augmentation changes the relative productivity of human and AI suppliers, the products on which AI constrains human prices, and the selection of supplied products. It therefore changes the automation loss itself. Multiplying the augmentation gain by an automation loss calculated for the original human distribution generally misses these interactions.

A special case illustrates when multiplication does suffice. If a common productivity factor scales both human and AI technologies, their relative productivities, selection, and markups remain unchanged. Normalized demand scales by that factor to the power $1-\sigma$, while consumption scales by its power $(\sigma-1)/(\sigma-2)$. In that case the proportional automation loss is unchanged. Augmentation of human producers alone need not have this property, as the next example shows.

D.2 An explicit threshold in the two-outcome example

Return to the examples in Section 4, with $\kappa_o = 0$, a successful human draw of productivity one with probability s, and a low draw $0 \leq z_L < 1$ with probability $1-s$. Competitive AI supplies only the low-outcome products, with fixed productivity $z_A > z_L$. Augmentation raises the productivity of exceptional human outcomes from one to $a \geq 1$, leaving the low human outcome and AI productivity unchanged. This includes the success-failure example when $z_L = 0$.

Relative labor demand gives AI's share of production labor as

$$\lambda(a) := \frac{(1-s)\mu^\sigma z_A^{\sigma-1}}{sa^{\sigma-1} + (1-s)\mu^\sigma z_A^{\sigma-1}}. \tag{58}$$

Stronger augmentation raises the productivity of the surviving human producers and reduces AI's labor share. It therefore changes both the returns supporting development and the extent to which unequal markups draw labor toward AI production.

Proposition **8** (Augmentation sufficient to restore consumption)
In the two-outcome economy, consumption relative to the baseline without AI is

$$\frac{C_A(a)}{C_b} = \left[\frac{sa^{\sigma-1}}{s + (1-s)z_L^{\sigma-1}}\left(1 - \frac{\lambda(a)}{\sigma}\right)\right]^{1/(\sigma-2)}\left[1 - \frac{(\rho-\delta)\lambda(a)}{\sigma\rho - \delta}\right]. \tag{59}$$

This ratio is strictly increasing in a. There is a unique threshold $a_C > 1$ at which consumption returns to its no-AI level. Define

$$a_0 := \left[\frac{s + (1-s)z_L^{\sigma-1}}{s}\right]^{1/(\sigma-1)}$$

as the augmentation needed to replace the lost ordinary contribution to expected profits at baseline demand. The consumption threshold satisfies

$$a_0 < a_C < a_0\left\{\mu\left[\frac{\sigma\rho - \delta}{(\sigma-1)\rho}\right]^{\sigma-2}\right\}^{1/(\sigma-1)}. \tag{60}$$

In the patient limit, these bounds simplify to $a_0 < a_C < \mu^{1/(\sigma-1)}a_0$. At $a = a_C$, develop-

ment remains below its no-AI level.

Proof. At baseline demand, expected human profits are proportional to $s + (1-s)z_L^{\sigma-1}$. With both augmentation and automation, only exceptional outcomes reward their developers, and their productivity moment is $sa^{\sigma-1}$. Free entry therefore gives

$$\frac{X_b}{X_A(a)} = \frac{sa^{\sigma-1}}{s + (1-s)z_L^{\sigma-1}}.$$

The labor-weighted average markup is $\bar{m}_A = \mu[1 - \lambda(a)/\sigma]$. Substitution into (55), with zero operating costs, yields (59). As a rises, $a^{\sigma-1}$ increases strictly and $\lambda(a)$ falls strictly. The consumption ratio consequently increases strictly, is below one at $a = 1$, and tends to infinity as $a \to \infty$. This establishes the unique threshold.

At $a = a_0$, the expected-profit moment is restored, but both markup-related factors are at most one and the first is strictly below one. Thus $a_C > a_0$. For the upper bound, use $0 < \lambda(a) < 1$ to obtain

$$1 - \frac{\lambda(a)}{\sigma} > \frac{1}{\mu}, \qquad 1 - \frac{(\rho-\delta)\lambda(a)}{\sigma\rho - \delta} \geq \frac{(\sigma-1)\rho}{\sigma\rho - \delta}.$$

The first inequality is strict, including in the patient limit. At the upper bound in (60), these inequalities make consumption strictly exceed C_b.

Finally, aggregate free entry gives $I_A/I_b = K_A/K_b = (1-\lambda)P_b/P_A$. Evaluating this expression where $C_A/C_b = 1$ yields

$$\left.\frac{I_A}{I_b}\right|_{a=a_C} = \frac{1 - \lambda(a_C)}{1 - \dfrac{(\rho-\delta)\lambda(a_C)}{\sigma\rho - \delta}} < 1.$$

When $z_L > 0$, every developed idea is supplied in both economies, so the same ratio applies to product variety. When $z_L = 0$, $M_b = sK_b$ and $M_A = K_A$; the distinction between development and supplied products remains necessary. □

The threshold a_0 isolates the loss of ordinary commercial returns emphasized in the introduction. Exceptional outcomes must become sufficiently profitable to replace this contribution to the expected payoff from development. Competitive automation requires still stronger augmentation because it also draws labor toward products with lower markups. Yet, recovering baseline consumption does not require recovering baseline development: more productive human producers and lower development expenditure can compensate for fewer ideas being developed. The exact threshold accounts for the fact that stronger augmentation itself reduces AI's labor share, rather than treating the automation loss as fixed.

E Product Development with Labor

The main text assumes that developing a product requires purchases of the composite good. We now consider an alternative in which product development employs workers directly. This changes the resource trade-off: development competes with production for labor, rather than with consumption for goods. It also removes the direct effect of lower goods prices on the cost of development.

Within this appendix, $\kappa_d > 0$ denotes the labor required to develop a product. We change only the development technology; the operating cost $\kappa_o \geq 0$ remains in units of the composite good. All other assumptions are unchanged. In particular, the wage is normalized to one, products fall out of fashion with probability δ, and $\rho = 1 - \beta(1-\delta)$. We continue to compare stationary equilibria and retain $\sigma > 2$.

E.1 The Baseline without AI

Development and production. Let L^p denote labor employed in production. Since developing I new products requires $\kappa_d I$ workers, labor market clearing and the goods resource constraint become

$$L^p := L - \kappa_d I, \qquad Y \equiv C + \kappa_o M, \qquad I = \delta K. \tag{61}$$

Development uses no goods directly, but it lowers the labor available to produce them. With zero operating costs, gross output therefore equals consumption.

As before, firms charge $p(z) = \mu/z$, where $\mu = \sigma/(\sigma-1)$. Let s^* denote the fraction of developed products that are supplied, so that $M = s^* K$, and let $\tilde{z}$ denote average supplier productivity, defined as in (4). Common markups imply

$$Y = M^{1/(\sigma-1)} \tilde{z} L^p, \qquad P^{-1} = \frac{M^{1/(\sigma-1)} \tilde{z}}{\mu}. \tag{62}$$

Thus, the real wage continues to measure productivity in goods production, while consumption also depends on how much labor is used to develop products.

Selection and entry. The individual operating profit remains $\pi(z) = (\mu-1)\ell(z) - P\kappa_o$, measured in wage units. However, free entry now equates the expected value of these profits to a labor cost:

$$\rho \kappa_d = s^* \pi(\tilde{z}). \tag{63}$$

For $\kappa_o > 0$ and an interior production cutoff z^*, the marginal producer earns zero profits. Relative labor demand then gives

$$\pi(z) = P\kappa_o \left[\left(\frac{z}{z^*} \right)^{\sigma-1} - 1 \right], \qquad P\kappa_o s^* \left[\left(\frac{\tilde{z}}{z^*} \right)^{\sigma-1} - 1 \right] = \rho\kappa_d. \tag{64}$$

Unlike in the main text, the price index no longer cancels from the selection condition. Operating expenses are paid in goods, while development costs are paid in labor. Consequently, the cutoff and the price index must be determined jointly.

Multiplying free entry by the stock of developed products and adding back operating expenses gives aggregate production labor:

$$(\mu - 1)L^p = K\big(\rho\kappa_d + P\kappa_o s^*\big).$$

Combining this condition with labor market clearing yields

$$K = \frac{L}{\kappa_d[\delta + (\sigma-1)\rho] + (\sigma-1)P\kappa_o s^*}, \qquad M = s^* K, \qquad I = \delta K. \tag{65}$$

Together with (62) and (64), these equations characterize an interior stationary equilibrium. Finally, the goods resource constraint and aggregate free entry imply

$$C = \frac{\mu L^p}{P} - \kappa_o M = \frac{L + (\rho - \delta)\kappa_d K}{P}. \tag{66}$$

As in the main text, consumption per worker converges to the real wage in the patient limit. Away from that limit, profits net of replacement development also contribute to consumption.

Zero operating costs. The comparison becomes particularly transparent when $\kappa_o = 0$, as in the two-outcome examples in Section 4. All products with positive productivity are supplied. Suppose $0 < \int_0^\infty z^{\sigma-1}\, dG(z) < \infty$. Equations (65) and (61) give

$$K_b = \frac{L}{\kappa_d[\delta + (\sigma-1)\rho]}, \qquad L_b^p = \frac{(\sigma-1)\rho}{\delta + (\sigma-1)\rho} L. \tag{67}$$

The stock of developed products and the allocation of labor are independent of the productivity distribution. A higher probability of obtaining a productive design increases supplied variety, but does not increase the number of development attempts. In particular,

$$P_b^{-1} = \frac{1}{\mu} \left[K_b \int_0^\infty z^{\sigma-1}\, dG(z) \right]^{1/(\sigma-1)}, \qquad C_b = Y_b = \frac{\mu L_b^p}{P_b}. \tag{68}$$

The reason is that aggregate monopoly profits equal $(\mu - 1)L_b^p$. Free entry equates these profits to $\rho\kappa_d K_b$, while labor market clearing allocates the remaining workers to replacement development. Higher productivity raises the purchasing power of both wages and profits, without changing their ratio to the wage cost of development.

E.2 AI Augmentation

As in Section 3, augmentation transforms an original draw u into effective productivity $z = \alpha(u)u$, with $\alpha(u) \geq 1$ and $\alpha'(u)u + \alpha(u) > 0$. The transformed distribution satisfies $G_\alpha(\alpha(u)u) \equiv G(u)$. Replacing G by G_α in the preceding equilibrium conditions gives the augmented economy, including the determination of selection when operating costs are positive.

PROPOSITION **9** (Augmentation with development costs in labor)
Suppose $\kappa_o = 0$ and augmentation preserves whether a draw has positive productivity. Then augmentation leaves development investment, supplied product variety, and production labor unchanged. The real wage, output, and consumption increase by the same factor:

$$\frac{P_b}{P_\alpha} = \frac{Y_\alpha}{Y_b} = \frac{C_\alpha}{C_b} = \left[\frac{\int_0^\infty [\alpha(u)u]^{\sigma-1}\, dG(u)}{\int_0^\infty u^{\sigma-1}\, dG(u)}\right]^{1/(\sigma-1)}. \tag{69}$$

Under uniform augmentation, $\alpha(u) = A$, this factor equals A.

Proof. Equation (67) determines K and L^p independently of the productivity distribution. Since augmentation preserves positive productivity, the supplied fraction is unchanged as well. Hence $I = \delta K$ and $M = s^* K$ remain unchanged. Substituting the transformed productivity moment into (68) gives (69). Uniform augmentation multiplies this moment by $A^{\sigma-1}$. □

Compared with Proposition 1, the additional gain through greater product variety disappears. Uniform augmentation raises production productivity, but does not make the labor required for development cheaper relative to the labor employed by producing firms. Nonuniform augmentation still matters through its effect on average supplier productivity: improvements to already productive firms receive greater weight in the CES productivity index. However, with zero operating costs, they do not change aggregate development investment.

With positive operating costs, augmentation also reduces the labor needed to produce the goods used for operation. Uniform augmentation can be interpreted exactly as a reduction in operating requirements in the unaugmented economy, followed by a proportional increase in output. To make this comparison explicit, write $K_b(\kappa_o)$ and $C_b(\kappa_o)$ for the baseline outcomes at operating requirement κ_o. For corresponding stationary

equilibria,

$$K_\alpha = K_b(\kappa_o/A), \qquad P_\alpha = \frac{P_b(\kappa_o/A)}{A}, \qquad C_\alpha = A\, C_b(\kappa_o/A). \tag{70}$$

Indeed, at a given allocation of labor, augmenting productivity by A multiplies output by A and divides prices by A. The monetary operating cost is then the same as in the unaugmented economy with operating requirement κ_o/A, while the labor cost of development is unchanged. Selection and development can therefore change when $\kappa_o > 0$, but through operating costs rather than through cheaper development goods.

E.3 Competitive Automation of Ordinary Success

We now return to the two viable outcomes in Section 4.3. Operating costs are zero. A developed product has human productivity one with probability $s \in (0,1)$ and productivity $z_L \in (0,1)$ otherwise. Competitive AI supplies all ordinary products at productivity $z_A \geq z_L$ and price $1/z_A$, without compensating their developers. At equality, AI serves the market. Exceptional human products remain protected and charge μ. All developed products are supplied, so $M = K$ in both economies.

Labor demand and development. Let λ denote AI's share of production labor, excluding workers employed in development. Relative demand gives the same expression as in the main text:

$$\lambda := \frac{(1-s)\mu^\sigma z_A^{\sigma-1}}{s + (1-s)\mu^\sigma z_A^{\sigma-1}}. \tag{71}$$

The difference is that production labor is now endogenous. Only exceptional products generate development returns, so aggregate free entry and labor market clearing imply

$$\rho\kappa_d K_A = (\mu-1)(1-\lambda)L_A^p, \qquad L_A^p = L - \delta\kappa_d K_A. \tag{72}$$

Consequently,

$$K_A = \frac{(1-\lambda)L}{\kappa_d[(\sigma-1)\rho + \delta(1-\lambda)]}, \qquad L_A^p = \frac{(\sigma-1)\rho L}{(\sigma-1)\rho + \delta(1-\lambda)}. \tag{73}$$

Automation lowers development and releases workers into production. The following proposition shows that this reallocation does not overturn the stationary consumption loss.

PROPOSITION **10** (Competitive automation with development costs in labor)
In the two-viable-outcome economy with competitive AI, product variety and development investment fall, while production labor increases:

$$\frac{M_A}{M_b} = \frac{I_A}{I_b} = \frac{1-\lambda}{1-\dfrac{\delta\lambda}{(\sigma-1)\rho+\delta}} < 1, \qquad \frac{L_A^p}{L_b^p} = \frac{1}{1-\dfrac{\delta\lambda}{(\sigma-1)\rho+\delta}} > 1. \tag{74}$$

Nevertheless, the real wage and consumption decline:

$$\frac{P_b}{P_A} = \left[\frac{s}{s+(1-s)z_L^{\sigma-1}}\frac{1-\lambda/\sigma}{1-\dfrac{\delta\lambda}{(\sigma-1)\rho+\delta}}\right]^{1/(\sigma-1)} < 1, \tag{75}$$

$$\frac{C_A}{C_b} = \underbrace{\left[\frac{s}{s+(1-s)z_L^{\sigma-1}}\right]^{1/(\sigma-1)}}_{\text{loss of ordinary development returns}} \underbrace{\left[\frac{1-\lambda/\sigma}{1-\dfrac{\delta\lambda}{(\sigma-1)\rho+\delta}}\right]^{\sigma/(\sigma-1)}}_{\text{competition and the reallocation of labor}} < 1. \tag{76}$$

In the patient limit, the second factor in (76) equals one. Both the consumption ratio and the real-wage ratio then equal $[s/(s+(1-s)z_L^{\sigma-1})]^{1/(\sigma-1)} < 1$.

Proof. Comparing (73) with (67) gives the variety and labor ratios. The two price indices satisfy

$$P_b^{-1} = \frac{[K_b(s+(1-s)z_L^{\sigma-1})]^{1/(\sigma-1)}}{\mu}, \qquad P_A^{-1} = \frac{[K_A(s+(1-s)(\mu z_A)^{\sigma-1})]^{1/(\sigma-1)}}{\mu}.$$

As in Section 4.3, the definition of AI's labor share implies

$$s+(1-s)(\mu z_A)^{\sigma-1} = \frac{s}{1-\lambda}\left(1-\frac{\lambda}{\sigma}\right).$$

Substituting the variety ratio into the price comparison gives (75). With zero operating costs, all goods are consumed. Output aggregation therefore gives

$$\frac{C_A}{C_b} = \left(1-\frac{\lambda}{\sigma}\right)\frac{L_A^p}{L_b^p}\frac{P_b}{P_A},$$

which yields (76). Finally, $\rho \geq \delta$ implies $\delta/[(\sigma-1)\rho+\delta] \leq 1/\sigma$. The second factor is thus at most one, and equals one when $\rho = \delta$. The first factor is strictly below one because ordinary outcomes previously generated positive development returns. □

The first factor again measures the effect of losing the ordinary contribution to expected development profits. The second combines the larger weight of competitive AI

production with the workers released from development. Its numerator, $1 - \lambda/\sigma$, is the same output wedge at a given real wage and production labor as in the main text. Its denominator captures the increase in production labor, which also supports development returns and hence product variety. With patient households these effects offset each other, but the loss of ordinary returns remains. For $\beta < 1$, the release of development workers is not large enough to offset the effects of competition. Thus, the negative stationary consumption result does not require development costs to be paid in goods, although their units change its magnitude and the way resources are reallocated.

F Transition after Competitive Automation

The stationary comparisons in Section 4.3 allow product development and the stock of varieties to adjust fully. Following an unexpected introduction of AI, however, the economy initially retains products developed under the previous incentives. We now derive the transition in the two-viable-outcome model and distinguish the change in stationary consumption from the welfare effect of introducing AI when these products already exist.

F.1 Timing and Equilibrium

We return to the main text's technology: development costs κ_d are paid in goods, operating costs are zero, and all labor is used in production. Human productivity is one with probability $s \in (0, 1)$ and $z_L \in (0, 1)$ otherwise. The economy is initially in the stationary equilibrium without AI, with stock K_b and consumption C_b. At the beginning of period zero, before new development and production, competitive AI becomes available unexpectedly and permanently. It immediately supplies existing and newly developed ordinary products at productivity $z_A \geq z_L$, with ties assigned to AI. Exceptional human products remain protected. Developers receive no compensation for automated products, and there is no further uncertainty about aggregate conditions.

Let $K_{-1} = K_b$ denote the stock used in production in the preceding period. Products fall out of fashion at the end of each period, as in the main text. Development $I_t \geq 0$ takes place before production and adds designs that can be supplied immediately. The stock used in period t therefore satisfies

$$K_t = (1 - \delta)K_{t-1} + I_t. \tag{77}$$

All remaining designs are supplied, so $M_t = K_t$. The continuum of draws and productivity-independent obsolescence keep the fraction of exceptional products equal to s in every cohort. Hence the stock of designs is the only aggregate state needed for the transition.

Production and consumption. AI's share of production labor is constant after its introduction:

$$\lambda = \frac{(1-s)\mu^{\sigma} z_A^{\sigma-1}}{s + (1-s)\mu^{\sigma} z_A^{\sigma-1}}.$$

Given the stock of products, the price index, output, and consumption are

$$P_t^{-1} = \frac{\left[K_t(s + (1-s)(\mu z_A)^{\sigma-1})\right]^{1/(\sigma-1)}}{\mu}, \tag{78}$$

$$Y_t = \mu\left(1 - \frac{\lambda}{\sigma}\right)\frac{L}{P_t}, \qquad C_t = Y_t - \kappa_d I_t. \tag{79}$$

At a given stock, AI affects production immediately. Development responds to future profits, while the stock of products can shrink only as existing products fall out of fashion.

Development incentives. Let V_t denote the expected value of a newly developed design, measured in goods, before its productivity is known. Let π_t^H denote the operating profit in wage units of an exceptional human product. Only this outcome rewards its developer. Since utility is linear in consumption, the continuation value satisfies

$$V_t = \frac{s\pi_t^H}{P_t} + \beta(1-\delta)V_{t+1}, \qquad V_t \le \kappa_d, \qquad I_t \ge 0, \qquad I_t(\kappa_d - V_t) = 0. \tag{80}$$

Values are the discounted sums of the corresponding profit flows. Free entry need not hold with equality during the transition: when developing a product costs more than its expected value, no new products are developed. Also, profits must be converted into goods period by period. Dividing a stationary profit flow by ρ would not account for the changing price index.

Let K_A, Y_A, and C_A denote the stationary stock, output, and consumption after automation. Aggregate human profits equal $(\mu-1)(1-\lambda)L$. Combining this fact with (78) and stationary free entry gives

$$\frac{s\pi_t^H}{P_t} = \frac{(\mu-1)(1-\lambda)L}{K_t P_t} = \rho\kappa_d\left(\frac{K_t}{K_A}\right)^{-(\sigma-2)/(\sigma-1)}. \tag{81}$$

Expected profits per design decrease with the stock of supplied products. When the inherited stock exceeds K_A, the market is too crowded to support continued development at its previous rate.

PROPOSITION **11** (The transition after competitive automation)
Following the unexpected introduction of competitive AI, the unique perfect-foresight transition with values equal to discounted profits satisfies

$$K_t = \max\{(1-\delta)^{t+1}K_b, K_A\}, \qquad I_t = \max\{K_A - (1-\delta)K_{t-1}, 0\}. \tag{82}$$

Define T as the first period in which the stock reaches its new stationary level:

$$T := \min\{t \geq 0 : (1-\delta)^{t+1}K_b \leq K_A\}. \tag{83}$$

For $t < T$, no products are developed and the stock declines through obsolescence. In period T, development is $I_T = K_A - (1-\delta)^{T+1}K_b$; from period $T+1$ onward it equals δK_A. Prices, output, and consumption follow (78)–(79). Consumption weakly decreases after its impact response and converges to $C_A < C_b$.

Proof. Section 4.3 establishes $K_A < K_b$. First, an equilibrium stock cannot fall below K_A. If $K_t < K_A$, (81) implies $s\pi_t^H/P_t > \rho\kappa_d$. Since $V_t \leq \kappa_d$, the value recursion then requires $V_{t+1} < \kappa_d$, so $I_{t+1} = 0$. The stock falls further and expected current profits increase. Repeating this argument would eventually make current profits alone exceed κ_d, contradicting $V_t \leq \kappa_d$ and nonnegative continuation profits.

Therefore $K_t \geq K_A$ at every date, and expected current profits are at most $\rho\kappa_d$. If $K_t > K_A$, the inequality is strict in the current period, implying $V_t < \kappa_d$ and $I_t = 0$. Positive development can occur only at $K_t = K_A$, where equality of value and cost also requires the stock to remain at K_A thereafter. Combining these restrictions with (77) gives (82), and the displayed path satisfies all entry conditions.

Before T, output falls with K_t and consumption equals output. At T, the stock reaches K_A and nonnegative development expenditure resumes. Since $0 \leq I_T \leq \delta K_A$, consumption in that period is at least C_A. From $T+1$ onward, the economy is stationary. This proves the consumption statement, including the case $T = 0$. □

The transition separates two effects that occur together in a stationary comparison. The introduction of AI changes the production technology and prices of ordinary goods immediately, while initially drawing on designs whose development costs have already been incurred. At the same time, lower development saves goods for consumption. The loss of product variety emerges over time, as fewer new designs replace those that fall out of fashion. Consequently, an initial consumption gain can coexist with the stationary consumption loss derived in the main text. The adjustment speed is governed by obsolescence and by the decline in the stationary stock, rather than by an assumed adjustment cost for development.

F.2 Welfare along the Transition

Let $W_A := \sum_{t=0}^{\infty} \beta^t C_t$ denote welfare following the unexpected introduction of AI, and let $W_b := C_b/(1-\beta)$ denote welfare if the economy remains in its initial stationary equilibrium. The comparison treats the development costs of inherited designs as sunk. The explicit stock path also gives an explicit welfare comparison:

$$W_A = Y_A \left(\frac{(1-\delta)K_b}{K_A} \right)^{1/(\sigma-1)} \frac{1 - \left[\beta(1-\delta)^{1/(\sigma-1)}\right]^T}{1 - \beta(1-\delta)^{1/(\sigma-1)}} + \frac{\beta^T C_A}{1-\beta} + \beta^T \kappa_d (\delta K_A - I_T). \tag{84}$$

The first term sums consumption during the periods without development. The second is the value of stationary consumption from period T onward. The last term accounts for the fact that reaching K_A in period T generally requires less development than subsequently maintaining it. These are resource effects; changes in the value of existing firms are not counted as a separate welfare loss.

Proposition **12** (Stationary losses and transition welfare)
In the two-viable-outcome model with competitive AI:

(i) If $z_A = z_L$, introducing AI lowers welfare, $W_A < W_b$.

(ii) For any fixed $\beta < 1$, sufficiently high AI productivity raises welfare, $W_A > W_b$, even though stationary consumption remains below C_b.

(iii) For any fixed, finite $z_A \geq z_L$, introducing AI lowers welfare when households are sufficiently patient. In particular,

$$\lim_{\beta \to 1} \frac{W_A}{W_b} = \left[\frac{s}{s + (1-s) z_L^{\sigma-1}} \left(1 - \frac{\lambda}{\sigma} \right) \right]^{1/(\sigma-2)} < 1. \tag{85}$$

Proof. Before period T, consumption equals $Y_A(K_t/K_A)^{1/(\sigma-1)}$. Summing this geometric sequence gives the first term of (84). In period T, consumption equals $C_A + \kappa_d(\delta K_A - I_T)$, and it equals C_A thereafter. This establishes the welfare expression.

For part (i), suppose $z_A = z_L$. At any given stock of designs, the technologies are unchanged. Production with common markups attains the highest output from these technologies and the available labor, while competitive AI introduces unequal markups. Write

$$Y_b(K) := L\left[K(s + (1-s) z_L^{\sigma-1})\right]^{1/(\sigma-1)}$$

for output with the original technologies and efficient labor allocation. Then $Y_t \leq Y_b(K_t)$. Since $\sigma > 2$, this output function is concave in K. Baseline free entry implies $Y_b(K_b)/(\sigma K_b) = \rho\kappa_d$, and hence $Y_b'(K_b) = \mu\rho\kappa_d$. Define $\Delta K_t := K_t - K_b \leq 0$, with

$\Delta K_{-1} = 0$. The stock law gives

$$\sum_{t=0}^{\infty} \beta^t (I_t - \delta K_b) = \rho \sum_{t=0}^{\infty} \beta^t \Delta K_t.$$

Consequently, concavity and the goods resource constraint imply

$$W_A - W_b \leq \left[Y_b'(K_b) - \rho\kappa_d\right] \sum_{t=0}^{\infty} \beta^t \Delta K_t = (\mu - 1)\rho\kappa_d \sum_{t=0}^{\infty} \beta^t \Delta K_t < 0.$$

Saved development goods do not compensate for the loss of output when AI provides no technological improvement.

For part (ii), as $z_A \to \infty$, the stationary stock K_A tends to zero, while the inherited stock $(1-\delta)K_b$ remains positive. There is therefore no development in period zero for sufficiently high z_A. Since $\lambda \to 1$, equations (78)–(79) imply

$$\lim_{z_A \to \infty} \frac{C_0}{z_A} = L\left[(1-\delta)K_b(1-s)\right]^{1/(\sigma-1)} > 0.$$

Thus C_0, and hence W_A, becomes arbitrarily large, while W_b is independent of AI productivity. All subsequent consumption remains nonnegative along the equilibrium path.

For part (iii), hold AI productivity fixed and let $\beta \to 1$. Both stationary stocks converge to strictly positive limits, so the number of transition periods remains bounded. Multiplying (84) by $1-\beta$ makes the finite transition terms vanish and leaves stationary consumption. Applying the stationary consumption comparison with $\rho \to \delta$ gives (85). □

These results qualify the interpretation of the stationary consumption losses. With no productivity improvement, automation lowers welfare even after accounting for saved development costs. With sufficiently productive AI, however, the gain from applying the new technology to inherited designs can dominate the subsequent losses. Greater patience gives more weight to the economy with fewer designs and less weight to this initial gain. The limits therefore differ: making AI arbitrarily productive at a fixed discount factor need not have the same implication as making households arbitrarily patient at a fixed AI productivity.

An illustration. Consider $\sigma = 4$, $s = 0.2$, $z_L = 0.5$, $\beta = 0.96$, and $\delta = 0.10$, and normalize $L = \kappa_d = 1$. Ordinary products then employ one third of production labor before automation. These parameters illustrate the transition rather than constitute a calibration. Table 1 reports three levels of AI productivity. The welfare change is the constant proportional change in baseline consumption that gives the same discounted utility, $(1-\beta)W_A/C_b - 1$.

Table 1: Stationary consumption and welfare after competitive automation

AI productivity	Impact C_0	Stationary C_A	Welfare	T
	Change relative to the no-AI baseline (%)			
$z_A = z_L = 0.5$	+13.6	−28.6	−19.1	11
$z_A = \mu z_L = 2/3$	+29.5	−31.5	−14.4	17
$z_A = 1$	+73.8	−33.8	+4.8	27

Notes. The introduction of AI is unexpected and occurs before period-zero development and production. T is the first period at the new stationary stock; development is zero in periods $0, \ldots, T-1$. Impact and stationary consumption changes are $100(C_0/C_b-1)$ and $100(C_A/C_b-1)$. Welfare is $100[(1-\beta)W_A/C_b-1]$.

In all three cases, the immediate reduction in development makes more goods available for consumption. When AI is no more productive than the ordinary human supplier, this initial saving cannot compensate for the subsequent loss of products and the less efficient allocation of production labor. Raising AI productivity to one instead generates a large enough initial production gain to increase discounted welfare by about 4.8%, although stationary consumption falls by about 33.8%. The example therefore illustrates a distinction between a lower stationary standard of living and the welfare effect of introducing a productive technology into an economy with an inherited stock of designs.