Rethinking Investments in Human Capital in the Age of Artificial Intelligence

Aleandro Palazzo

May 2026

Abstract

How does Artificial Intelligence reshape human capital investment decisions? I develop an Overlapping Generations model where risk-averse agents choose among different training paths to maximize their lifetime utility. Agents form expectations about the future impact of AI in each training path, but some paths are more uncertain than others. Using US data from BLS employment projections and recent AI exposure metrics, the model reveals three key results. First, uncertainty is an important driver of reallocation: risk-averse individuals systematically avoid highly uncertain training paths (reducing employment by up to 7% in the most uncertain one), which endogenously increases expected wages in those training paths due to labor scarcity. Second, workers reallocate toward paths with higher expected demand and lower uncertainty; however, the resulting crowding effect dampens the positive demand shock, ultimately leaving expected wages substantially unchanged. Third, due to these general equilibrium forces the net impact of AI on expected wages remains modest across all training paths, contained between -6% and +6% compared to pre-AI (2022) median wages.

1 Introduction

For most of the twentieth century, technological progress has pointed in the direction of incentivizing investments in human capital. Successive waves of automation primarily displaced routine, mechanical tasks concentrated in low and middle-skill occupations, while simultaneously raising the demand for high-skill workers with higher education (Acemoglu and Autor, 2011; Autor et al., 2003). In this environment, a college degree has been a realtively safe investment, offering high expected earnings and low employment risk. The secular rise in the college wage premium and in educational attainment throughout the latter half of the twentieth century is largely a product of this dynamic (Autor, 2019).

Artificial Intelligence (AI) fundamentally disrupts this long-standing relationship. Unlike prior automation technologies, modern AI systems demonstrate a high degree of proficiency precisely in the non-routine, cognitive tasks that have historically been the exclusive domain of high-skill workers (Eloundou et al., 2024; Webb, 2019; Tomlinson et al., 2025; Massenkoff and McCrory, 2026). The implications are potentially far-reaching: the very occupations that attracted generations of students to invest heavily in human capital are now among the most exposed to this new technology.

A critical feature of this technological breakthrough, however, is its profound *uncertainty.* The net effect of AI on labor demand is not only unknown, but its expected impact and the magnitude of the uncertainty vary dramatically across occupations. The nature of this uncertainty is due to the ambigous impact of any technology which can substitute human labor at the task level. As standard in a task-based model, when some tasks are automated, two opposing effect arises: the substitution effect, that acts as a negative force by depressing labor demand for the groups of workers directly exposed (Acemoglu et al., 2025) and a productivity effect, that raises demand for all the other tasks not automated in the economy (Acemoglu and Restrepo, 2018; Gans and Goldfarb, 2026). Which of this two effects will eventually dominate in each occupation is unknown, but people form expectations based on their current information and they take decision regarding their investments in human capital already today. The key point of this paper is that, individuals do not simply choose based on expectations about the future, but they also take into account the uncertainty surrounding those expectations. For example, jobs that have a relative low exposition to AI (e.g. technical manual jobs) have a lower uncertainty about the future, probably the substitution effect for them will be low and they will benefit from the productivity effect. On the other hand, jobs that have a high exposition to AI (e.g. legal jobs) have a high uncertainty about the future, many tasks in their occupations could be automated and we do not know if the productivity effect will dominate or not.

To formalize and quantify these dynamics, I develop an Overlapping Generations (OLG) model in which heterogeneous, risk-averse agents choose among J training paths to maximize their discounted expected utility of future wages, net of idiosyncratic investment costs. The key innovation is the decomposition of the AI shock for each training path j into two distinct components: an expected productivity shift $\mu_{AI,j}$, capturing the expected change in labor demand; and a variance component $\sigma^2_{AI,j}$, capturing the uncertainty about AI's ultimate impact in each training path. Agents are risk averse and invest in human capital to maximize their expected utility. The individual discrete choice problem aggregates to a tractable multinomial logit equilibrium.

The model delivers three main theoretical results, which I prove formally and then quantify through calibration to US data: An increase in AI-driven uncertainty unambiguously reduces equilibrium employment in the exposed training path. This is intuitive and it is driven by risk averse individuals. The same increase in uncertainty raises the expected equilibrium wage in the exposed path. This is a pure general equilibrium effect: the flight of risk-averse workers reduces labor supply, creating labor scarcity that, given the future expected labor demand, pushes expected wages upwards. This endogenous risk premium arises independently of any direct productivity gain it is entirely driven by the labor supply response to uncertainty. Third, When AI raises expected productivity in a training path, it attracts new workers. This inflow of labor creates congestion that erodes marginal productivity and dampens individual wage gains: the increase in expected wages is strictly less than the productivity shift, and the difference grows with the labor supply elasticity across training paths.

These three mechanisms interact in ways that are difficult to detect using reduced-form methods alone. A training path experiencing high AI uncertainty may see its expected wage rise due to the endogenous risk premium, even if AI has a mildly negative expected productivity impact. Conversely, a training path receiving very positive AI news may see disappointing wage growth due to labor market congestion. This paper provides a structural framework to disentangle these effects.

I calibrate the model to the US labor market using data from two primary sources. The expected AI shock, at the occupational level, is measured using the change in 10-year BLS occupational employment projections between the pre-AI (2021) and post-AI (2024) projection windows. The uncertainty parameter is derived from the gap between the theoretical exposure of an occupation to AI (share of tasks performable by AI; Eloundou et al. 2024) and its current actual exposure (Massenkoff and McCrory, 2026). The employment shares across 14 training paths are constructed from the IPUMS American Community Survey (ACS) for 2022, while the median wages by occupation are collected from the Bureau of Labor Statistics (BLS) for 2022. I recover the structural training cost parameters by inverting the pre-AI steady-state equilibrium (2022), and simulate the resulting post-AI equilibrium.

The quantitative results confirm the importance of the uncertainty channel: for several training paths, the uncertainty effect on employment is comparable in magnitude to the expected productivity shift, and in some cases reverses its sign entirely. For instance, under the baseline scenario, architecture—the most uncertain path—sees a net employment decline of 4% relative to its pre-AI share despite a positive expected productivity shock, as risk-averse agents flee the high uncertainty. On the wage side, the crowding effect substantially dampens expected wage gains in training paths experiencing positive expected productivity shocks, while for training paths with high uncertainty their expected wages rise significantly due to the endogenous risk premium. Overall, because these general equilibrium forces (productivity shifts and labor supply adjustments) partially offset one another, the net impact of AI on expected wages remains modest across all training paths, contained between -6% and +6%.

This paper contributes to three strands of the literature. First, it builds on the large empirical literature on AI's impact on the labor market (among many others, Acemoglu and Restrepo 2018; Webb 2019; Agrawal et al. 2025; Ide and Talamàs 2025; Acemoglu et al. 2025; Gans and Goldfarb 2026), by providing a structural model that allows for a clean decomposition of the supply-side response to the demand and uncertainty channels. Second, it contributes to the literature on human capital investment under uncertainty (Levhari and Weiss, 1974; Cunha et al., 2005; Cunha and Heckman, 2016), by introducing a new form of risk in the context of human capital investment decisions. Third, it connects to the macro literature on skill-biased technological change (Acemoglu and Autor, 2011; Autor et al., 2003; Autor, 2019) by characterizing the novel mechanism through which AI reshapes economic incentives in ways that are qualitatively different from prior automation waves.

The rest of the paper is organized as follows. Section 2 develops the theoretical model. Section 3 describes the measurement of the AI shock. Theoretical results and equilibrium properties are derived in Section 2.4. Section 4 describes the calibration strategy, and Section 5 presents the quantitative simulation results. Finally, Section 6 concludes.

2 The Model

2.1 Environment

I consider an Overlapping Generations (OLG) economy where agents live for two periods: young ($g = y$) and old ($g = o$). The economy features a finite set of potential training paths. Each training path involves a period of training (when young) followed by a period of trained employment (when old). The only exception is the case of no training, which represents a fallback option with no investment in human capital. Training paths may be of two types: a *technical* training path (e.g., plumber, cook) is charachterized by a period of apprenticeship as training, while a *cognitive* training path involves a period of studying in college in a certain field before entering the labor market. Each field of study (in the cognitive training path) is associated with many occupations, whereas a *technical* training path is associated with one occupation. Let's denote the different training paths by $j \in \{0, 1, ..., J\}$ ($j = 0$ refers to the case of no training). When young ($g = y$), agents are endowed with a unit of time and choose a training path j. This choice involves an investment in human capital with uncertain returns. In the second period ($g = o$), agents supply labor inelastically, realize their wage (which depends on the chosen training path and labor market conditions), and consume.

2.2 Demand Side

A representative firm demands labor from each training path j according to the production function:

$$Y = \sum_{j=0}^{J} A_j L_j^{\alpha} e^{\theta_j}$$

where L_j is the share of working population with training j, A_j is TFP for training path j, and θ_j is an AI-induced labor demand shock specific to training path j. In particular, $\theta_j = 0$ pre-AI, while it becomes stochastic post-AI and follows a normal distribution with training path-specific mean and variance: $\theta_j \sim \mathcal{N}(\mu_{AI,j}, \sigma^2_{AI,j})$. The demand side is very simple on purpose, as the main focus of the paper is on the supply-side response to the AI shock. However, from the worker's perspective, it is not too distant from reality: they observe the current productivity level of each training path, A_j, the share of workers in that path, L_j, and they have some beliefs about the future impact of AI, which are captured by the distribution of θ_j. Intuitively, if $E[\theta_j] = \mu_{AI,j} > 0$ AI is expected to increase labor demand for training path j, while if $\mu_{AI,j} < 0$ AI is expected to decrease labor demand. On the other hand, $Var[\theta_j] = \sigma^2_{AI,j}$ captures the uncertainty about the future impact of AI on labor demand in training path j.
Assuming competitive markets, the wage for labor with training j is:

$$\frac{\partial Y_j}{\partial L_j} = \alpha A_j e^{\theta_j} L_j^{\alpha-1}$$

Taking logs, the log-wage is:

$$w_j = \ln \alpha + \ln A_j + \theta_j - (1-\alpha) \ln L_j \tag{1}$$

Clearly, the wage depends on the (unknown) AI shock θ_j and on the employment level L_j, which is endogenously determined by the agents' training choices.

2.3 Supply Side

Agents are risk-averse, with preferences described by a Constant Absolute Risk Aversion (CARA) utility function over the log-wage $w = \ln W$: $u(w) = -e^{-\rho w}$.[1] Individuals

[1]This is mathematically equivalent to assuming a Constant Relative Risk Aversion (CRRA) utility function defined over the wage level W, $u(W) = -\frac{W^{1-\gamma}}{1-\gamma}$, where $\rho = \gamma - 1$ is the coefficient of absolute risk aversion.

are heterogeneous with respect to a vector of net costs $C_i = [C_{i0}, C_{i1}, \ldots, C_{iJ}]$, which aggregates all non-pecuniary disutility and direct expenses specific to each training path. For example, if a worker really hates math, the cost of choosing engineering will be huge, whereas if she really loves literature, the cost of choosing humanities will be very low.

A young agent i at time t selects the training path j that maximizes the discounted expected utility of future earnings, net of the investment cost:

$$V_{it}^{y}(C_i) = \max_j \left\{\beta \mathbb{E}_t \left[u(w_{j,t+1})\right] - C_{ij}\right\} \tag{2}$$

where $\beta \in (0,1)$ is the discount factor and the log-wage for training path j is given by the marginal product of labor:

$$w_{j,t} = \ln \alpha + \ln A_j + \theta_{j,t} - (1-\alpha) \ln L_{j,t}$$

Because the AI productivity shock θ_j is distributed as $\mathcal{N}(\mu_{AI,j}, \sigma_{AI,j}^2)$, the log-wage w_j is normal. Hence, the expected utility can be evaluated exactly using the moment-generating function of the normal distribution:

$$\mathbb{E}_t[u(w_j)] = \mathbb{E}_t[-e^{-\rho w_j}] = -\exp\left(-\rho \mathbb{E}_t[w_j] + \frac{\rho^2}{2}\text{Var}_t(w_j)\right)$$

The certainty equivalent Ω_j is defined as the guaranteed wage that yields the same expected utility, $u(\Omega_j) = \mathbb{E}_t[u(w_j)]$. Equating the utility of Ω_j to the exact expected utility yields:

$$-e^{-\rho \Omega_j} = -\exp\left(-\rho \mathbb{E}_t[w_j] + \frac{\rho^2}{2}\text{Var}_t(w_j)\right)$$

Taking logs and rearranging provides the exact certainty equivalent for training path j:

$$\Omega_j = \mathbb{E}_t[w_j] - \frac{\rho}{2}\text{Var}_t(w_j) = \underbrace{\ln \alpha + \ln A_j + \mu_{AI,j} - (1-\alpha)\ln L_j}_{\mathbb{E}[w_j]} - \underbrace{\frac{\rho}{2}\sigma_{AI,j}^2}_{\text{Risk premium}}$$

which formalizes exactly how agents penalize uncertainty.

Substituting $\mathbb{E}_t[u(w_j)]$ with $u(\Omega_j)$, the agent's objective function becomes:

$$V_{it}^{y}(C_i) = \max_j \left\{\beta u(\Omega_j) - C_{ij}\right\}$$

To render the discrete choice problem analytically tractable, I apply a first-order Taylor expansion to the utility of the certainty equivalent $u(\Omega_j)$ around a baseline average wage $\bar{w}$. Substituting $u(\Omega_j) \approx u(\bar{w}) + u'(\bar{w})(\Omega_j - \bar{w})$ into the objective, the maximization problem simplifies to:

$$\max_j \left\{ \beta \Omega_j - C^*_{ij} \right\} \tag{3}$$

where $C^*_{ij} \equiv C_{ij}/u'(\bar{w})$ is the normalized cost parameter.

I assume this normalized cost is determined by a training-specific common average cost, $\bar{C}^*_j$, and an idiosyncratic preference component, such that $C^*_{ij} = \bar{C}^*_j - \epsilon_{ij}$. Assuming the idiosyncratic shocks ϵ_{ij} are independently and identically distributed following a Type I Extreme Value (Gumbel) distribution with location parameter zero and scale parameter 1, the aggregate market share of labor choosing training path j is mapped by the standard multinomial conditional logit formulation (McFadden (1972)):

$$L^*_j = \frac{\exp\left(\beta \Omega_j(L^*_j) - \bar{C}^*_j\right)}{\sum_{h=0}^{J} \exp\left(\beta \Omega_h(L^*_h) - \bar{C}^*_h\right)} \tag{4}$$

2.4 Theoretical Results and Equilibrium Properties

Using the structural discrete choice framework derived above, this section establishes the comparative statics of the model with respect to the two key moments of the AI shock: the expected labor demand shift ($\mu_{AI,j}$) and the variance of the shock ($\sigma^2_{AI,j}$). All formal mathematical proofs are relegated to the Appendix.

Proposition 1 (Flight from Uncertainty)**.** *An increase in AI-driven uncertainty for training path j ($\uparrow \sigma^2_{AI,j}$) leads to a decrease in the equilibrium employment L^*_j:*

$$\frac{\partial L^*_j}{\partial \sigma^2_{AI,j}} = -\frac{\frac{\rho\beta}{2} L^*_j(1 - L^*_j)}{1 + \beta(1-\alpha)} < 0 \tag{5}$$

Intuition: As a career path becomes more unpredictable due to exposure to AI, risk-averse individuals systematically flee to safer alternatives. The magnitude of this structural flight depends directly on their degree of risk aversion (ρ).

Proposition 2 (Endogenous Risk Premium)**.** *An increase in AI-driven uncertainty for training path j ($\uparrow \sigma^2_{AI,j}$) raises the expected wage $\mathbb{E}[w_j]^*$. This reflects an endogenous risk premium driven by labor scarcity:*

$$\frac{\partial \mathbb{E}[w_j]^*}{\partial \sigma^2_{AI,j}} = -\frac{1-\alpha}{L^*_j} \frac{\partial L^*_j}{\partial \sigma^2_{AI,j}} > 0 \tag{6}$$

Intuition: This characterizes one of the core theoretical mechanisms of the model. Uncertainty reduce the share of labor choosing training path j (see proposition 1).

Keeping fixed the expected labor demand, the reduction in labor supply leads to an increase in the expected wage.

Proposition 3 (AI Productivity and the Crowding Effect). *An increase in the expected future labor demand for training path j ($\uparrow \mu_{AI,j}$) leads to an increase in both the equilibrium employment L_j^* and the expected wage $\mathbb{E}[w_j]^*$. However, the increase in the expected wage is less than proportional:*

$$\frac{\partial L_j^*}{\partial \mu_{AI,j}} = \frac{\beta L_j^*(1 - L_j^*)}{1 + \beta(1 - \alpha)} > 0 \tag{7}$$

$$0 < \frac{\partial \mathbb{E}[w_j]^*}{\partial \mu_{AI,j}} = \frac{1 + \beta(1 - \alpha)L_j^*}{1 + \beta(1 - \alpha)} < 1 \tag{8}$$

Intuition: If AI is expected to act as a broad complementary technology for training path j ($\mu_{AI,j} > 0$), it shifts the expected marginal product outward. While this raises expected wages, it structurally triggers a inflow of new workers. This labor supply response creates congestion that erodes the marginal product and dampens expected wages.

Proposition 4 (Cross-Career Spillovers - Employment). *An increase in the expected future labor demand ($\uparrow \mu_{AI,j}$) or a decrease in AI-driven uncertainty for training path j ($\downarrow \sigma^2_{AI,j}$) leads to a decrease in the employment level of all other careers $h \neq j$:*

$$\frac{\partial L_h^*}{\partial \mu_{AI,j}} = -\frac{\beta L_h^* L_j^*}{1 + \beta(1 - \alpha)} < 0 \tag{9}$$

$$\frac{\partial L_h^*}{\partial \sigma^2_{AI,j}} = \frac{\frac{\beta\rho}{2} L_h^* L_j^*}{1 + \beta(1 - \alpha)} > 0 \tag{10}$$

Intuition: Human capital supply is a zero-sum game. If AI makes training path j more attractive, it inevitably drains workers from other paths h. Conversely, if AI makes career j less attractive, workers actively seek refuge in relatively positively affected careers.

Proposition 5 (Cross-Career Spillovers - Expected Wages). *An increase in the expected future labor demand ($\uparrow \mu_{AI,j}$) or a decrease in AI-driven uncertainty for training path j ($\downarrow \sigma^2_{AI,j}$) leads to an increase in the expected wages of all other careers $h \neq j$:*

$$\frac{\partial \mathbb{E}[w_h]^*}{\partial \mu_{AI,j}} = -(1 - \alpha)\frac{1}{L_h^*}\frac{\partial L_h^*}{\partial \mu_{AI,j}} > 0 \tag{11}$$

$$\frac{\partial \mathbb{E}[w_h]^*}{\partial \sigma^2_{AI,j}} = -(1 - \alpha)\frac{1}{L_h^*}\frac{\partial L_h^*}{\partial \sigma^2_{AI,j}} < 0 \tag{12}$$

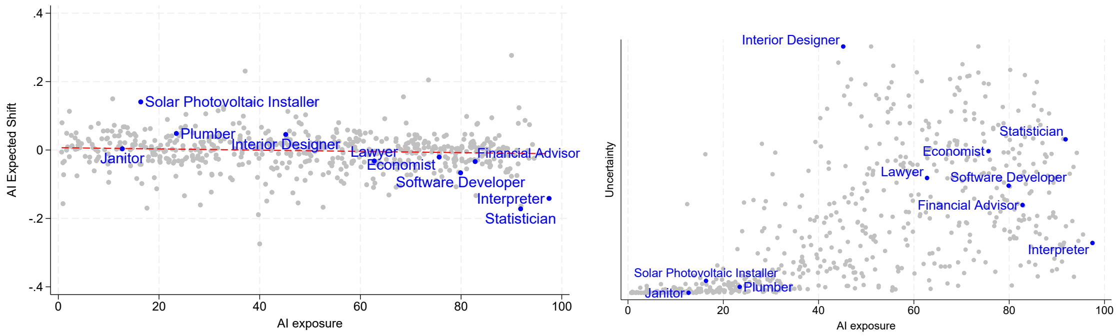

Figure 1: AI-induced expected labor demand shift ($\mu_{AI,k}$) and uncertainty ($\sigma^2_{AI,k}$) across occupations.

Intuition: These cross-career employment shifts induce an inverse, endogenous general equilibrium wage dynamic. If an overwhelming share of individuals rushes into safer paths (less risky or more appealing), the emerging shortage of workers in the other paths mechanically increases their expected wages.

3 Measuring the AI Shock

To empirically measure the first two moments of the AI shock θ_j for each training path j, I propose a two-step approach. First, I estimate the same two moments for each occupation, and then I aggregate these occupation-level estimates to the training path level using a weighted average.

Hence, suppose that the impact of AI on occupation k is given by $\theta_k \sim \mathcal{N}(\mu_{AI,k}, \sigma^2_{AI,k})$. To estimate the expected labor demand shift ($\mu_{AI,k}$), I use the 10-year employment projected growth rates provided by the Bureau of Labor Statistics (BLS) for each occupation k. Specifically, $\mu_{AI,k}$ is calculated as the percentage change between the 2024 projected growth (post-AI) and the 2021 projected growth (pre-AI). This is used as the baseline measure of the expected shift.

To estimate the uncertainty parameter ($\sigma^2_{AI,k}$), I exploit the variation in the gap between the *theoretical exposure* (T_k) and the *actual exposure* (A_k). Theoretical exposure T_k represents the share of tasks within occupation k that could theoretically be performed by AI (according to Eloundou et al. (2024)), and actual exposure A_k indicates the share of tasks currently being performed by AI (according to Massenkoff and McCrory (2026)). I define the uncertainty (or the variance) of the AI shock for occupation k as

$$\sigma^2_{AI,k} = \left(\frac{T_k - A_k}{2}\right)^2 \tag{13}$$

Finally, I map these occupational shocks to specific training paths j. The path-level

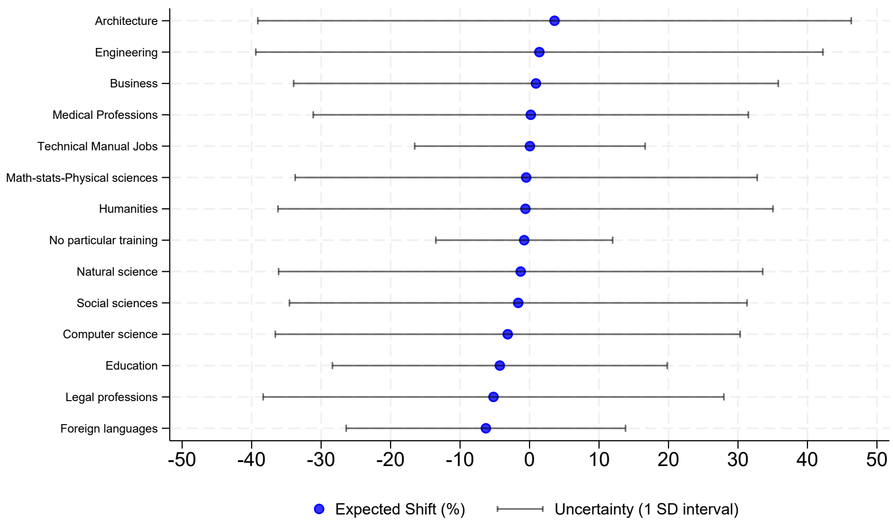

Figure 2: AI-induced expected labor demand shift ($\mu_{AI,j}$) and uncertainty ($\sigma^2_{AI,j}$) across training paths.

mean and variance are obtained via an employment-weighted average of the constituent occupations:

$$\mu_{AI,j} = \sum_{k \in j} \omega_{j,k} \mu_{AI,k} \quad \text{and} \quad \sigma^2_{AI,j} = \lambda \sum_{k \in j} \omega_{j,k} \sigma^2_{AI,k} \tag{14}$$

where $\omega_{j,k}$ is the share of workers with training j employed in occupation k. The scaling factor λ is needed to make the variance parameter comparable to the expected labor demand shift, and it is calibrated in Section 4 in order to match a reasonable value for the risk premium.

Figure 1 shows the distribution of the expected labor demand shift and the uncertainty across occupations, while Figure 2 shows the resulting distribution across training paths after the aggregation step. Notice that the expected labor demand shift across occupations ranges between -20% and +20%. Also the uncertainty is quite heterogeneous across occupations, with some of them (e.g., architecture, engineering) showing a very large uncertainty, and others (e.g., technical manual jobs) showing a much smaller standard deviation. This heterogeneity is crucial for the results of the paper, as individuals take into account both the expected labor demand shift and the relative uncertainty of each training path.

4 Calibration

To quantify the structural impact of AI on the US labor market, I calibrate the parameters of the model to match the cross-sectional distribution of wages and employment observed in the pre-AI steady state (2022). This allows me to recover the latent costs of human capital investment and baseline productivity levels for each of the $J = 14$ training paths.

4.1 Parameter Selection

The core structural parameters are calibrated to the standard macro literature. The labor share in the production function is set to $\alpha = 0.65$, and the discount factor is $\beta = 0.8$, reflecting the two-period structure of the OLG model. I assume a coefficient of absolute risk aversion $\rho = 2$ (equivalent to a relative risk aversion $\gamma = 3$).

Table 1: Model Calibration Summary

Parameter	**Symbol**	**Value / Source**	**Target / Rationale**
Externally set parameters			
Labor share	α	0.65	Standard
Discount factor	β	0.8	$\approx$ 0.99 yearly
Absolute risk aversion	ρ	2.0	Relative risk aversion = 3
Structurally inverted from pre-AI data (2022)			
Training costs	$\bar{C}_j^*$	Cross-section	Logit inversion (ACS 2022 shares & wages)
Baseline TFP	$\ln \bar{A}_j$	Cross-section	Marginal product condition (ACS 2022)
AI shock moments (occupation $\rightarrow$ training path aggregation)			
Expected AI shift	$\mu_{AI,j}$	BLS projections	Δ 10-yr growth rate (2024 vs. 2021)
AI uncertainty	$\sigma^2_{AI,j}$	Task exposure gap	$\left(\frac{T_k - A_k}{2}\right)^2$, rescaled
Variance scaling	λ	Estimated under 3 scenarios	Risk premium scenarios (5%, 10%, 15%)

Notes: ACS = IPUMS American Community Survey. BLS = Bureau of Labor Statistics. T_k = theoretical AI task exposure (Eloundou et al., 2024); A_k = actual AI task exposure (Massenkoff and McCrory, 2026). All monetary variables are in log units.

I define the "No Particular Training" category as the reference training path ($j = 0$). I normalize its investment cost to $C_0^* = 0$. Given the observed pre-AI conditional logit shares L_j^{pre} and median wages w_j^{pre} from the IPUMS American Community Survey (ACS, 2022), I invert the optimal choice condition to recover the structural cost parameters:

$$C_j^* = \beta(\ln w_j^{pre} - \ln w_0^{pre}) - \ln \left(\frac{L_j^{pre}}{L_0^{pre}} \right) \tag{15}$$

Consistent with the demand-side production function, the latent baseline productivity level ($\bar{A}_j$) for each path is recovered from the pre-AI marginal product condition:

$$\ln \bar{A}_j = \ln w_j^{pre} - \ln \alpha + (1 - \alpha) \ln L_j^{pre} \tag{16}$$

Finally, the variance is scaled in order to be comparable to the expected labor demand shift, and to match a reasonable value for the risk premium. Specifically, I choose λ such that the average risk premium across all paths under the baseline scenario is equal to 10%. Then, I perform a sensitivity analysis by scaling λ to match risk premium scenarios of 5% (conservative) and 15% (high risk premium).
Table 1 summarises all calibrated parameters.

5 Results and Decomposition

This section presents the quantitative predictions of the calibrated model. To isolate the contribution of each channel, I use a sequential decomposition. The two outcome variables — employment and expected wages — are decomposed differently, reflecting the distinct economic mechanisms at play:

- **Employment decomposition:** The *no uncertainty allocation* is the employment shift that would occur if AI only changed expected productivity ($\mu_{AI,j} \neq 0$, $\sigma^2_{AI,j} = 0$), i.e., in the absence of any uncertainty. The *uncertainty effect* is the residual: the additional employment reallocation driven by risk-averse agents fleeing more volatile training paths.

- **Wage decomposition:** The *demand effect* is the wage change under the post-AI productivity parameters but holding the labor distribution fixed at its pre-AI baseline — isolating the pure direct technology impact on expected wages with no reallocation. The *crowding effect* is the residual: the general-equilibrium wage adjustment caused by workers actually moving across sectors in response to the full shock.

5.1 Employment Effects

Figure 3 presents the employment decomposition across the 14 training paths. Following the sequential decomposition, I first isolate the *no uncertainty allocation* — the employment shift that would occur if AI only affected expected labor demand with no uncertainty ($\sigma^2_{AI,j} = 0$). Technical manual jobs, medical professions, business, architecture, and engineering would gain workers given their positive expected productivity shocks, while the other training paths would lose workers. I then add the *uncertainty effect* (blue diamonds), which captures the additional reallocation driven by risk aversion: agents actively avoid training paths with high variance (hence more uncertain), even when the expected wage remains relatively attractive. The total effect (black squares) reveals that the uncertainty channel is quantitatively significant, either reversing the sign or amplifying the no uncertainty reallocation. For example architecture, the most uncertain path, would see a net employment shift of -4% even though the expected productivity shock is positive.

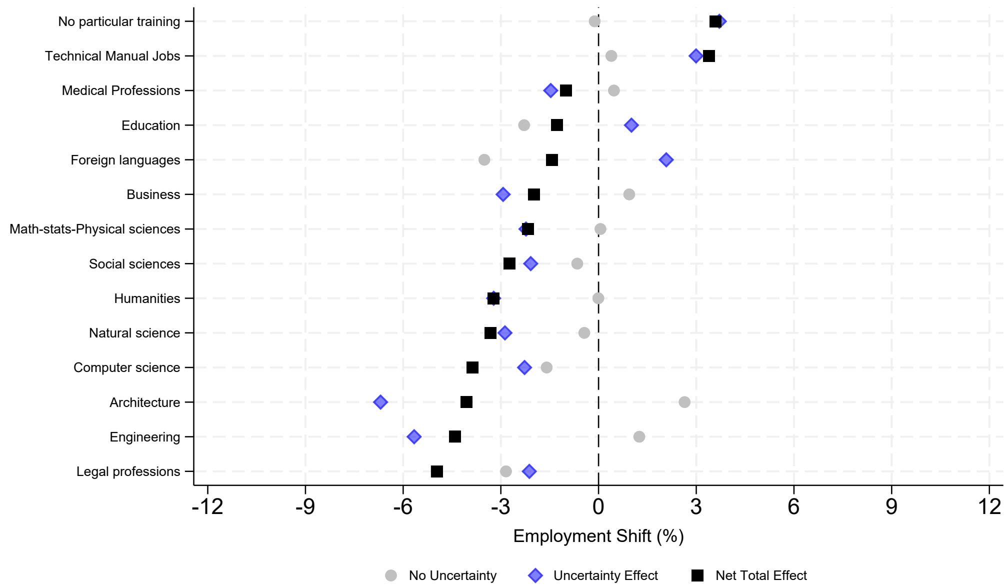

Figure 3: Decomposition of AI-induced employment shifts across training paths. Grey circles: no uncertainty allocation ($\sigma^2_{AI} = 0$). Blue diamonds: uncertainty effect (risk-aversion induced reallocation). Black squares: total net effect. Employment shifts expressed as percentage change relative to the pre-AI baseline share.

5.2 Wage Effects

Figure 4 presents the decomposition of expected wage changes. I first compute the *demand effect* (grey dots): the direct wage impact of the AI productivity shock, computed by applying the post-AI parameters $\mu_{AI,j}$ to the *pre-AI* labor distribution — isolating the pure technology channel with no reallocation. I then add the *crowding effect* (blue diamonds), which captures the general equilibrium wage adjustment caused by workers actually moving across sectors in response to the full shock. Training paths attracting large inflows see downward wage pressure as labor congestion erodes marginal productivity; training paths losing workers experience upward wage pressure through the endogenous risk premium. Two patterns stand out. First, the crowding effect consistently *dampens* expected wages in relatively safer training paths. Second, for sectors experiencing a flight from uncertainty, the crowding effect is *positive*: labor scarcity pushes expected wages above what the pure productivity shock would predict. Crucially, these two forces partially offset one another, so that the net impact on expected wages is not drastic across all paths, contained between -6% and $+6\%$.

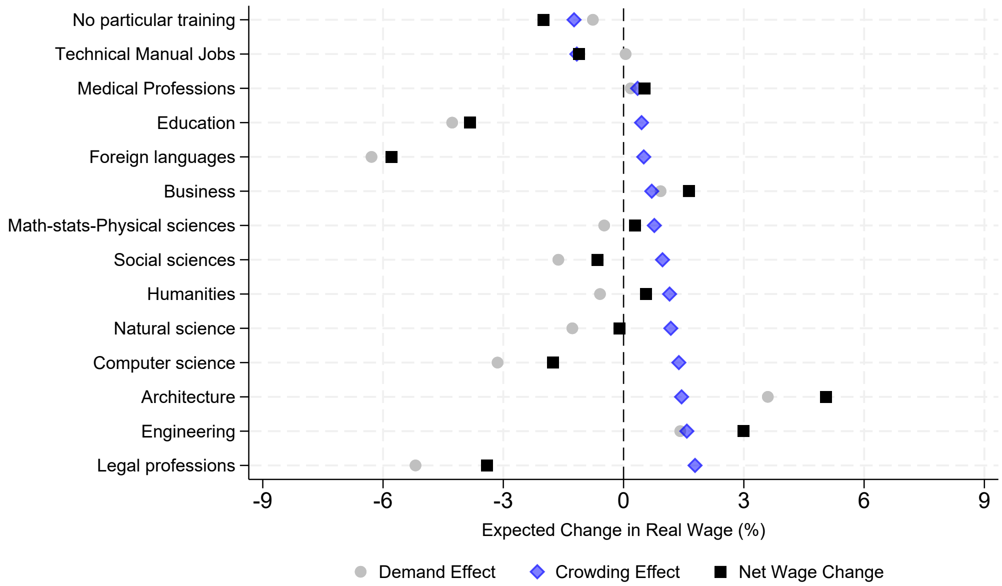

Figure 4: Decomposition of AI-induced expected wage changes across training paths. Grey circles: demand effect (post-AI productivity applied at pre-AI labor distribution). Blue diamonds: crowding effect (wage adjustment due to endogenous workers' reallocation). Black squares: total net wage change.

5.3 Sensitivity to Extreme Uncertainty

Given the irreversibility of human capital investments, the true risk premium demanded by students facing an unprecedented structural shock like AI could be significantly higher than historical benchmarks. Figures 5 and 6 present the sensitivity of the net employment shift and expected wage change across the three risk premium scenarios: 5% (Low), 10% (Baseline), and 15% (High).

As the risk premium increases, the uncertainty channel dominates the demand channel. For occupations highly exposed to uncertainty, the flight from uncertainty causes an accelerating collapse in employment shares (Figure 5). Concurrently, this extreme labor scarcity mechanically amplifies the endogenous risk premium, pushing expected wages significantly for those few who remain in the exposed careers (Figure 6). This highlights that if AI is perceived as an extreme tail risk, the reallocation of labor will be massive and primarily driven by risk avoidance. However, the net wage impact is relativley lower in magnitude than the employment shift, reflecting the fact that productivity and crowding effects partially offset one another.

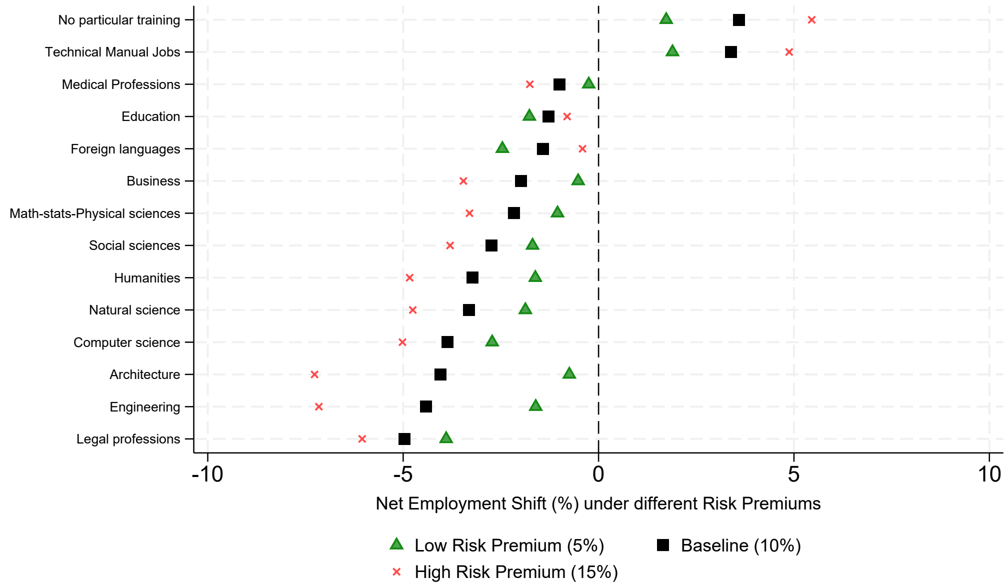

Figure 5: Sensitivity analysis: Net employment shift (%) under three different risk premium scenarios.

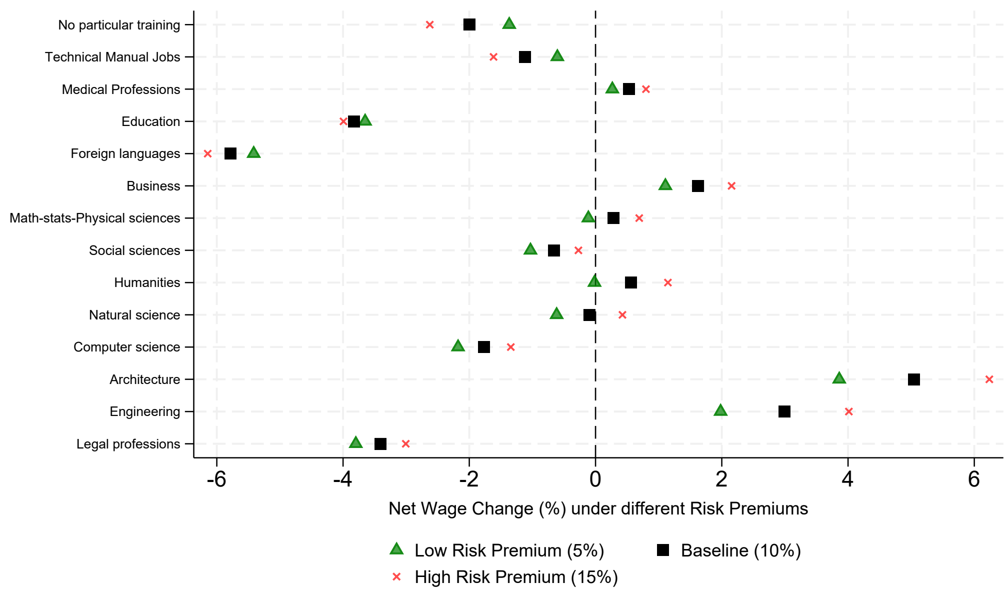

Figure 6: Sensitivity analysis: Net wage change (%) under three different risk premium scenarios.

6 Conclusion

This paper investigates how Artificial Intelligence reshapes human capital investment decisions when workers face technological uncertainty. To study this, I develop a structural Overlapping Generations (OLG) model in which risk-averse agents choose among different training paths, evaluating both the expected demand shifts and the variance surrounding AI's future impact. By calibrating the model to the US labor market using ACS data, task-level AI exposure metrics, and BLS occupational projections, I decompose the post-AI labor market equilibrium into direct demand, general equilibrium crowding, and risk-aversion-driven uncertainty channels.

The theoretical and quantitative findings highlight three main insights. First, technological uncertainty is a major determinant of career choices: risk-averse individuals systematically flee highly uncertain training paths (such as architecture and engineering), leading to a baseline employment drop of up to 4%—and up to 7% in the high risk premium sensitivity scenario—even when expected demand is positive. Second, this flight creates labor scarcity that drives up expected wages in exposed sectors, generating an endogenous risk premium that compensates the remaining workers. Third, the crowding of workers into perceived "safe havens" creates congestion that dampens expected wage gains. Overall, due to these offsetting general equilibrium forces, the net impact of AI on expected wages remains modest across all training paths, contained between -6% and +6%.

These findings have important implications for labor policy and educational planning. Traditional employment forecasts that only focus on point estimates of expected demand may overlook the powerful role of uncertainty in shaping labor supply.

Appendix

A Mathematical Appendix: Derivations of Equilibrium Dynamics

A.1 Derivation of the Employment Equilibrium

To derive the labor share choosing each training path j, I aggregate the discrete choices of heterogeneous agents. The derivation follows four logical steps.

A.1.1 Step 1: Individual Utility and Idiosyncratic Shocks

The utility of individual i choosing career j is defined as the discounted certainty equivalent of the future wage net of investment costs:

$$V_{ij} = \beta\Omega_j - C_{ij}$$

The cost parameter C_{ij} is decomposed into a career-specific mean $\bar{C}_j$ and an individual-specific idiosyncratic shock ϵ_{ij}:

$$V_{ij} = \underbrace{\beta\Omega_j - \bar{C}_j}_{U_j} + \epsilon_{ij}$$

where U_j represents the average utility of career j. Let's assume that the shocks ϵ_{ij} are i.i.d. following a Gumbel distribution (Extreme Value Type I).

A.1.2 Step 2: The Decision Rule

An agent chooses career j if it provides the maximum utility among all $J+1$ alternatives:

$$V_{ij} \geq V_{ih} \implies U_j + \epsilon_{ij} \geq U_h + \epsilon_{ih} \quad \forall h \in \{0, \ldots, J\}$$

This can be rewritten as a condition on the idiosyncratic shocks:

$$\epsilon_{ih} < \epsilon_{ij} + U_j - U_h \quad \forall h \neq j$$

A.1.3 Step 3: Aggregation via Integration

For a fixed value of ϵ_{ij}, the probability that career j is chosen is the product of the probabilities that all other ϵ_{ih} are below the threshold:

$$P(j|\epsilon_{ij}) = \prod_{h\neq j} F(\epsilon_{ij} + U_j - U_h) = \prod_{h\neq j} \exp\left(-e^{-(\epsilon_{ij}+U_j-U_h)}\right)$$

The aggregate labor supply L_j^* is found by integrating this conditional probability over the entire distribution of ϵ_{ij} using its PDF $f(\epsilon) = e^{-\epsilon}\exp(-e^{-\epsilon})$:

$$L_j^* = \int_{-\infty}^{\infty} \left[\prod_{h\neq j} \exp\left(-e^{-\epsilon_{ij}-U_j+U_h}\right)\right] e^{-\epsilon_{ij}} \exp(-e^{-\epsilon_{ij}}) d\epsilon_{ij}$$

Combining the exponential terms (including the case $h = j$ from the PDF):

$$L_j^* = \int_{-\infty}^{\infty} e^{-\epsilon_{ij}} \exp\left(-e^{-\epsilon_{ij}} \sum_{h=0}^{J} e^{-(U_j-U_h)}\right) d\epsilon_{ij}$$

A.1.4 Step 4: Change of Variables and Final Equilibrium

Let $t = e^{-\epsilon_{ij}}$, then $dt = -e^{-\epsilon_{ij}} d\epsilon_{ij}$. Changing the limits of integration and solving the integral:

$$L_j^* = \int_0^{\infty} \exp\left(-t \sum_{h=0}^{J} \frac{e^{U_h}}{e^{U_j}}\right) dt = \frac{1}{\sum_{h=0}^{J} \frac{e^{U_h}}{e^{U_j}}}$$

Rearranging gives the multinomial logit market share:

$$L_J^* = \frac{\exp(U_J)}{\sum_{h=0}^{J} \exp(U_h)} = \frac{\exp(\beta\Omega_J - \bar{C}_J^*)}{\sum_{h=0}^{J} \exp(\beta\Omega_h - \bar{C}_h^*)}$$

A.2 Derivation of Equation (A1)

Equation (A1) follows directly from the utility-maximization problem under Type I Extreme Value shocks. The derivation can be written in four steps.

Step 1: Define systematic utility. For each training path j, write the individual utility as

$$V_{ij} = U_j + \epsilon_{ij}, \qquad U_j \equiv \beta\Omega_j - \bar{C}_j^*$$

where U_j is the deterministic component and ϵ_{ij} is the idiosyncratic taste shock.

Step 2: State the choice event. Agent i chooses j if and only if

$$V_{ij} \geq V_{ih} \quad \forall h \neq j$$

which is equivalent to

$$\epsilon_{ih} \leq \epsilon_{ij} + U_j - U_h \quad \forall h \neq j.$$

Step 3: Integrate over the Gumbel distribution. Because the shocks are i.i.d. Type I Extreme Value, the conditional choice probability is

$$P(j \mid \epsilon_{ij}) = \prod_{h \neq j} F(\epsilon_{ij} + U_j - U_h),$$

and integrating over the density of ϵ_{ij} delivers the standard multinomial logit share.

Step 4: Recover the closed form. Thus,

$$L_j^* = \frac{\exp(U_j)}{\sum_{h=0}^{J} \exp(U_h)} = \frac{\exp(\beta\Omega_j - \bar{C}_j^*)}{\sum_{h=0}^{J} \exp(\beta\Omega_h - \bar{C}_h^*)} \tag{A1}$$

which is exactly equation (A1).

A.3 Derivation of the Logit Identity

Starting from the multinomial logit share,

$$L_j^* = \frac{\exp(U_j)}{\sum_{h=0}^{J} \exp(U_h)},$$

take natural logarithms on both sides:

$$\ln L_j^* = U_j - \ln\left(\sum_{h=0}^{J} \exp(U_h)\right).$$

Differentiating with respect to a generic parameter x gives

$$\frac{\partial \ln L_j^*}{\partial x} = \frac{\partial U_j}{\partial x} - \frac{1}{\sum_{h=0}^{J} \exp(U_h)} \sum_{h=0}^{J} \exp(U_h) \frac{\partial U_h}{\partial x}.$$

Using the definition of the logit share,

$$\frac{\exp(U_h)}{\sum_{m=0}^{J} \exp(U_m)} = L_h^*,$$

the expression simplifies to

$$\frac{\partial \ln L_j^*}{\partial x} = \frac{\partial U_j}{\partial x} - \sum_{h=0}^{J} L_h^* \frac{\partial U_h}{\partial x}, \tag{A2}$$

which is the desired identity.

A.4 Standard Logit Identities

Consider the market share for career j in a multinomial logit framework:

$$L_j^* = \frac{\exp(U_j)}{\sum_{h=0}^{J} \exp(U_h)}$$

Taking the natural logarithm and differentiating with respect to any parameter x, we obtain the fundamental identity for logit shares:

$$\frac{\partial \ln L_j^*}{\partial x} = \frac{\partial U_j}{\partial x} - \sum_{h=0}^{J} L_h^* \frac{\partial U_h}{\partial x} \tag{A3}$$

Furthermore, since $\sum_{h=0}^{J} L_h^* = 1$, it must hold that $\sum_{h=0}^{J} \frac{\partial L_h^*}{\partial x} = 0$.

A.5 Proof of Proposition 1 (Flight from Uncertainty)

The average utility for career j is given by $U_j = \beta\Omega_j - \bar{C}_j^*$, where Ω_j is the certainty equivalent:

$$U_j = \beta \left[\ln \alpha + \ln \bar{A}_j + \mu_{AI,j} - (1-\alpha) \ln L_j^* - \frac{\rho}{2} \sigma_{AI,j}^2 \right] - \bar{C}_j^*$$

Differentiating U_j and U_h $(h \neq j)$ with respect to $\sigma_{AI,j}^2$:

$$\frac{\partial U_j}{\partial \sigma_{AI,j}^2} = -\beta(1-\alpha) \frac{1}{L_j^*} \frac{\partial L_j^*}{\partial \sigma_{AI,j}^2} - \frac{\beta\rho}{2}$$
$$\frac{\partial U_h}{\partial \sigma_{AI,j}^2} = -\beta(1-\alpha) \frac{1}{L_h^*} \frac{\partial L_h^*}{\partial \sigma_{AI,j}^2} \quad (h \neq j)$$

Substituting these into identity (A3):

$$\frac{1}{L_j^*} \frac{\partial L_j^*}{\partial \sigma_{AI,j}^2} = \left[-\frac{\beta(1-\alpha)}{L_j^*} \frac{\partial L_j^*}{\partial \sigma_{AI,j}^2} - \frac{\beta\rho}{2} \right] - \sum_{h=0}^{J} L_h^* \frac{\partial U_h}{\partial \sigma_{AI,j}^2}$$

The summation term simplifies as follows:

$$\sum_{h=0}^{J} L_h^* \frac{\partial U_h}{\partial \sigma_{AI,j}^2} = L_j^* \left(-\frac{\beta\rho}{2} \right) - \beta(1-\alpha) \underbrace{\sum_{h=0}^{J} \frac{\partial L_h^*}{\partial \sigma_{AI,j}^2}}_{0} = -\frac{\beta\rho}{2} L_j^*$$

Rearranging the terms to isolate $\frac{\partial L_j^*}{\partial \sigma_{AI,j}^2}$:

$$\frac{\partial L_j^*}{\partial \sigma_{AI,j}^2}\left[\frac{1}{L_j^*}+\frac{\beta(1-\alpha)}{L_j^*}\right]=-\frac{\beta\rho}{2}+\frac{\beta\rho}{2}L_j^*$$

Multiplying by L_j^*:

$$\frac{\partial L_j^*}{\partial \sigma_{AI,j}^2}[1+\beta(1-\alpha)]=-\frac{\beta\rho}{2}L_j^*(1-L_j^*)$$

Which yields the final result:

$$\frac{\partial L_j^*}{\partial \sigma_{AI,j}^2}=-\frac{\frac{\beta\rho}{2}L_j^*(1-L_j^*)}{1+\beta(1-\alpha)}<0 \quad \blacksquare$$

A.6 Proof of Proposition 2 (Endogenous Risk Premium)

The expected log-wage is $\mathbb{E}[w_j]^*=\ln\alpha+\ln\bar{A}_j+\mu_{AI,j}-(1-\alpha)\ln L_j^*$. Differentiating with respect to $\sigma_{AI,j}^2$:

$$\frac{\partial \mathbb{E}[w_j]^*}{\partial \sigma_{AI,j}^2}=-\frac{1-\alpha}{L_j^*}\frac{\partial L_j^*}{\partial \sigma_{AI,j}^2}>0$$

since $\frac{\partial L_j^*}{\partial \sigma_{AI,j}^2}<0 \quad \blacksquare$

A.7 Proof of Proposition 3 (AI Productivity and Crowding)

Differentiating the log-share identity (A3) with respect to $\mu_{AI,j}$:

$$\frac{1}{L_j^*}\frac{\partial L_j^*}{\partial \mu_{AI,j}}=\beta\left[1-\frac{1-\alpha}{L_j^*}\frac{\partial L_j^*}{\partial \mu_{AI,j}}\right]-\beta L_j^*$$

Solving for $\frac{\partial L_j^*}{\partial \mu_{AI,j}}$:

$$\frac{\partial L_j^*}{\partial \mu_{AI,j}}=\frac{\beta L_j^*(1-L_j^*)}{1+\beta(1-\alpha)}>0$$

For the expected wage response:

$$\frac{\partial \mathbb{E}[w_j]^*}{\partial \mu_{AI,j}}=1-(1-\alpha)\frac{1}{L_j^*}\frac{\partial L_j^*}{\partial \mu_{AI,j}}=1-\frac{\beta(1-\alpha)(1-L_j^*)}{1+\beta(1-\alpha)}$$

Simplifying the fraction:

$$\frac{\partial \mathbb{E}[w_j]^*}{\partial \mu_{AI,j}}=\frac{1+\beta(1-\alpha)L_j^*}{1+\beta(1-\alpha)}\in(0,1) \quad \blacksquare$$

References

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, volume 4, pages 1043–1171. Elsevier.

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340.

Acemoglu, D., Kong, F., and Restrepo, P. (2025). Tasks at work: comparative advantage, technology and labor demand. *Handbook of Labor Economics*, 6:1–114.

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American economic review*, 108(6):1488–1542.

Agrawal, A. K., Gans, J. S., and Goldfarb, A. (2025). Genius on demand: The value of transformative artificial intelligence. Technical report, National Bureau of Economic Research.

Autor, D. H. (2019). Work of the past, work of the future. *AEA Papers and Proceedings*, 109:1–32.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics*, 118(4):1279–1333.

Cunha, F. and Heckman, J. (2016). Decomposing trends in inequality in earnings into forecastable and uncertain components. *Journal of labor economics*, 34(S2):S31–S65.

Cunha, F., Heckman, J., and Navarro, S. (2005). Separating uncertainty from heterogeneity in life cycle earnings. *oxford Economic papers*, 57(2):191–261.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). Gpts are gpts: Labor market impact potential of llms. *Science*, 384(6702):1306–1308.

Gans, J. S. and Goldfarb, A. (2026). O-ring automation. Technical report, National Bureau of Economic Research.

Ide, E. and Talamàs, E. (2025). Artificial intelligence in the knowledge economy.

Levhari, D. and Weiss, Y. (1974). The effect of risk on the investment in human capital. *The American Economic Review*, 64(6):950–963.

Massenkoff, M. and McCrory, P. (2026). Labor market impacts of ai: A new measure and early evidence.

McFadden, D. (1972). Conditional logit analysis of qualitative choice behavior.

Tomlinson, K., Jaffe, S., Wang, W., Counts, S., and Suri, S. (2025). Working with ai: measuring the applicability of generative ai to occupations. *arXiv preprint arXiv:2507.07935*.

Webb, M. (2019). The impact of artificial intelligence on the labor market. *Available at SSRN 3482150*.