



MASTERARBEIT | MASTER'S THESIS

Titel | Title

AI, Cognitive Inequality and Poverty Dynamics

verfasst von | submitted by

Emma Kramárová BSc (WU)

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of
Master of Science (MSc)

Wien | Vienna, 2026

Studienkennzahl lt. Studienblatt | Degree
programme code as it appears on the
student record sheet:

UA 066 913

Studienrichtung lt. Studienblatt | Degree
programme as it appears on the student
record sheet:

Masterstudium Applied Economics

Betreut von | Supervisor:

Assoz. Prof. Daniel Garcia PhD

Abstract- Deutsch

Künstliche Intelligenz soll die Produktivität steigern, doch die Voraussetzungen, um von ihr zu profitieren, sind weltweit ungleich verteilt. Die vorliegende Masterarbeit untersucht, ob die produktivitätswirksamen Erträge der KI-Nutzung im Ländervergleich vom kognitiven Kompetenzniveau der Erwerbsbevölkerung abhängen. Sie erweitert das task-basierte KI-Modell von *Acemoglu (2024)*, indem die beiden Parameter, die den aggregierten Produktivitätseffekt von KI steuern, als steigende Funktionen der kognitiven Kompetenz modelliert werden. Daraus folgt die überprüfbare Vorhersage, dass der Grenzertrag der KI-Bereitschaft mit den kognitiven Kompetenzen steigt. Empirisch umgesetzt wird das Modell als länderübergreifende OLS-Interaktionsspezifikation auf 98 Ländern, mit annualisierter TFP-Wachstumsrate als abhängiger Variable, dem Government AI Readiness Index als Maß für die KI-Bereitschaft und den Harmonised Learning Outcomes als Maß für kognitive Kompetenzen. Die Ergebnisse schließen einen Komplementärmechanismus zwischen KI und kognitiven Kompetenzen nicht aus, zeigen aber, dass er mit den verfügbaren Daten und beim gegenwärtigen Diffusionsstand im Länderquerschnitt nicht von null zu unterscheiden ist.

Abstract- English

Artificial intelligence is expected to raise productivity, but the conditions that allow economies to benefit from it are unequally distributed. This thesis asks whether the productivity returns to AI adoption depend on the cognitive skill level of the workforce. It extends Acemoglu's (2024) task-based model by treating the two parameters governing AI's aggregate productivity effect as increasing functions of cognitive skill, yielding the testable prediction that the marginal return to AI readiness rises with cognitive skills. The model is tested through a cross-country OLS interaction specification on 98 countries, using annualised TFP growth as the outcome, the Oxford Insights Government AI Readiness Index as the measure of AI readiness, and World Bank Harmonised Learning Outcomes as the measure of cognitive skills. The results do not rule out a skill-AI complementarity but indicate that, at the current diffusion stage and with available measurement, the mechanism cannot be distinguished from zero at the cross-country level.

TABLE OF CONTENTS

1. Introduction.....	5
1.1. Motivation.....	5
1.2. Research question.....	6
1.3. Contribution.....	6
2. Literature Review	6
2.1. AI, Automation and Task-based Technological Change	6
2.2. Human Capital, Cognitive Skills and Growth	7
2.3. Technology diffusion and development.....	8
2.4. Skill distribution, inequality and inclusive growth.....	10
2.5. Research gap.....	10
3. Theoretical framework.....	11
3.1. Baseline task-based AI Model	11
3.1.1. Production Structure	11
3.1.2. Four Channels of AI's Economic Impact.....	11
3.1.3. Hulten's Theorem and Aggregate Productivity.....	12
3.1.4. Quantitative Implications.....	12
3.1.5. Limitations of the Baseline Model and Motivation for Extension.....	12
3.2. Extending the model with cognitive skill level	13
3.2.1. Cognitive skill level as a modifier of AI productivity	13
3.2.1.1.....Introducing the cognitive skill parameter.....	13
3.2.1.2.....The skill utilization function.....	13
3.2.1.3.....Reinterpreting $\bar{\pi}$ and χ as country-specific	13
3.2.2. Formalizing the skill-adjusted productivity function.....	14
3.2.3. Testable Implications	14
4. Data and empirical strategy	15
4.1. Data sources.....	15
4.2. Measurement of AI Readiness R_c	16
4.2.1. Analytical variable	16
4.3. Measurement of cognitive skill level h_c	17
4.3.1. Data source and primary measure	17
4.3.2. Normalization and temporal Imputation	17
4.4. Outcome variables and controls	18
4.4.1. Primary Outcome: annualised TFP Growth	18
4.4.2. Conditioning variables	18
4.5. Empirical strategy	19
4.5.1. Cross-sectional baseline	19
4.5.2. Inference	20
4.5.3. Diagnostics	21
4.5.4. Sample construction and attrition.....	21
4.6. Robustness checks	21
5. Results.....	22
5.1. Sample description and descriptive statistics	22
5.2. Diagnostics.....	23
5.3. Baseline results M1 and M2 (H1 Test)	24
5.4. Marginal effects and threshold probing (H2).....	26
5.5. Distributional Channel H3 (M3, M3b)	27

5.6. Robustness (R4, R5, R16)	28
5.7. Within-country identification check (Panel FE)	28
6. Limitations.....	29
7. Discussion and Conclusion	29
8. References	30
Appendix.....	33

1. Introduction

1.1. Motivation

Artificial intelligence (AI) is widely seen as a transformative force for productivity, market creation, and economic change, yet the conditions that allow countries to benefit from AI are distributed highly unequally across the world. The World Bank's Digital Progress and Trends Report 2025 frames these conditions as the "4Cs" (connectivity, compute, context, and competency) and warns that AI innovation remains concentrated in high-income economies. By the second half of 2025, roughly one in six people worldwide used generative AI, with working-age adoption reaching 24.7% in the Global North but only 14.1% in the Global South. Left unaddressed, this divide risks turning AI into a force that widens global disparities. (*World Bank, 2025; Microsoft, 2026*)

AI has moved from commercial use into the public administration: the OECD documents governments deploying AI across internal operations, forecasting, fraud detection, accountability, and service delivery, and from 200 examined use scenarios identifies improved anomaly detection, better forecasting and decision-making, and more automated, customized procedures as the most anticipated benefits. The IMF reports that by mid-2023, at least 50 nations had set out AI strategies and 70 countries had launched AI-related policy initiatives. The reach and stakes of AI are thus expanding in parallel with the diffusion gap, raising the question of which countries can translate policy ambition into inclusive gains. (*Comunale & Manera, 2024; OECD, 2025*)

Nations differ significantly in their ability to adopt AI. The World Bank highlights obstacles in emerging economies: inadequate digital infrastructure, erratic connectivity, restricted device access, lack of affordable compute, and insufficient locally relevant data. These are fundamental requirements, determining whether AI can be adopted, adapted, and scaled (*World Bank, 2025*). Nations also vary in cognitive abilities and digital preparedness. *Hanushek and Woessmann (2008)* show that cognitive skills, rather than years of schooling, are the human-capital measure that matters for economic growth. The World Bank identifies "competency" (strong digital skills and wider capabilities) as essential for embracing, adapting, and inventing with AI (*World Bank, 2025; Hanushek & Woessmann, 2008*).

If AI's economic effects depend on foundational capabilities, and cognitive skills are a key determinant of long-run economic performance, then cross-country differences in cognitive skills may shape whether AI generates broad-based gains or reinforces existing inequalities. This is especially relevant since diffusion evidence shows adoption moving faster in advanced economies, while macroeconomic research remains cautious about large aggregate gains. *Acemoglu's* task-based framework shows that AI's productivity effects may be modest and that AI need not reduce inequality, it may even increase it or widen the gap between capital and labor income. The IMF similarly finds empirical data on employment and productivity remains ambiguous, despite theory predicting broad implications. (*Comunale & Manera, 2024; Microsoft, 2026; Acemoglu, 2024*)

This evidence suggests that AI's economic impacts are unlikely to depend only on technology diffusion, but also on whether nations possess the cognitive capacities to adopt and apply AI broadly.

1.2. Research question

The thesis asks the following central question: Does the impact of AI adoption on economic growth depend on the level of cognitive skills within a country?

To answer this question, the thesis investigates four related dimensions:

- Does the relationship between AI readiness and TFP growth vary with average cognitive skill levels?
- Is this relationship stronger in countries with higher average levels of cognitive proficiency?
- Are these effects robust to influential observations, sample composition, and an alternative outcome window?
- (Secondary) Is there evidence that the poverty-reducing effect of AI readiness is conditioned by cognitive skill levels?

1.3. Contribution

This thesis brings together two normally separate strands of literature: the emerging macroeconomics of AI and the literature on cognitive skills, human capital, and technology diffusion. Its central contribution is to test whether the economic returns to AI adoption vary systematically with cognitive skill levels across countries.

Theoretically, it extends *Acemoglu's (2024)* task-based framework by treating the two parameters governing AI's aggregate productivity effects (average task-level cost savings and the effective GDP share of tasks impacted by AI) as functions of the workforce's cognitive capabilities rather than fixed constants, allowing cross-country differences in foundational skills to condition AI's productivity effects.

Empirically, it combines cross-country indicators of AI adoption and readiness with internationally comparable measures of cognitive skills and macroeconomic outcomes. Total factor productivity growth is the primary outcome, with particular attention to the share of the population reaching basic cognitive proficiency as a theoretically motivated threshold.

2. Literature Review

2.1. AI, Automation and Task-based Technological Change

A significant portion of the economic literature views AI as the latest stage in a longer historical process of automation. Earlier technological waves such as steam power, electricity and software expanded the range of tasks performed by machines, while AI extends this process into non-routine and cognitive activities. *Aghion, Jones, and Jones* model AI as "the latest form of automation," highlighting that it may automate tasks, including higher-skill activities, previously believed to be beyond machines. According to the IMF's analysis, most economic research on AI uses task-based frameworks rather than aggregate human-capital categories to examine its implications on employment, wages, productivity, and growth (*Aghion et al., 2017; Comunale & Manera, 2024*).

The task-based approach builds on automation models in which output is produced through tasks allocated either to labor or capital. In the Zeira-style formulation discussed by *Aghion, Jones, and Jones*, non-automated tasks are produced by labor and automated tasks by capital (an increase in automated tasks raises the capital share and can increase long-run growth through capital accumulation). This simple framework predicts rising growth rates and capital shares as automation expands, conflicting with the relative stability of these aggregates over much of the twentieth century. *Aghion, Jones, and Jones* therefore adopt the Acemoglu-Restrepo strategy, which allows for automation of current jobs as well as creation of new ones, combining technological progress with steady development and factor shares by allowing automation and labor reinstatement to coexist (*Aghion et al., 2017*).

Acemoglu's (2024) model identifies four channels through which AI affects production (automation, task complementarity, deepening of automation, and new task creation). The framework shows that AI's aggregate productivity effects are disciplined by a version of Hulten's theorem, which leads *Acemoglu* to a more restrained conclusion than many public forecasts. *Saam (2024)* reaches similar conclusions, contending that the macroeconomic consequences of current AI diffusion, while noticeable, will not be particularly significant compared to previous digital changes (*Acemoglu, 2024; Saam, 2024*).

AI exposure differs from previous automation waves. According to the IMF analysis, white-collar, highly educated, and higher-paid occupations are significantly more exposed to AI, while robot-based automation was concentrated among blue-collar and less educated individuals. Clerical occupations have particularly high exposure, and augmentation effects may outweigh pure automation effects in both high- and low-income countries. The estimates indicate that 5.5% of all jobs in high-income nations are exposed to automation and 13.4% are augmented by AI. The comparable numbers for low-income countries are 0.4% and 10.4%. Whether AI will lower labor demand or rearrange work to increase productivity remains empirically unclear (*Comunale & Manera, 2024*).

The IMF review concludes that theoretical research broadly agrees AI will affect most occupations and transform growth, but empirical findings remain inconclusive on employment and productivity effects. *Acemoglu* contends that even if AI increases lower-skilled workers' productivity in specific jobs, inequality may still increase and the gap between capital and labor income may deepen. The literature suggests that AI should be analyzed neither as a purely labor-saving technology nor as an automatically inclusive productivity shock. Its effects depend on the interaction between the four channels described above and the structure of the economy into which it diffuses (*Acemoglu, 2024; Comunale & Manera, 2024*).

2.2. Human Capital, Cognitive Skills and Growth

A primary finding of the human capital literature is that educational quantity is not a sufficient indicator of abilities valuable to the economy. *Hanushek and Woessmann* argue that cross-country assessments based solely on years of schooling ignore the qualitative aspect of learning. The population's cognitive abilities, rather than educational attainment, are closely linked to economic development, individual earnings, and wealth distribution. Models incorporating direct measurements of cognitive skills explain significantly more variation in

long-term growth than models based only on years of schooling, and the coefficient on years of schooling becomes negligible once cognitive skills are included. This is especially important in international comparisons, where countries may have similar attainment but very different levels of learning (*Hanushek & Woessmann, 2008*).

The same analysis highlights the importance of cognitive skills across the distribution, not only on average. *Hanushek and Woessmann* question whether progress depends on a broad foundation of fundamental talents or on a small elite of high achievers. They employ two thresholds on the international test-score scale: 400 points for basic reading and 600 points for top performance and find separate relationships between economic growth and the share of pupils achieving basic proficiency and the share of top achievers. The effect of cognitive skills remains robust at different points of the distribution rather than only at the mean. They do not find a significant effect of cognitive-skill inequality per se on growth, but they do find significant effects at different points of the distribution, particularly at the basic proficiency threshold. This motivates the thesis's focus on average cognitive skill levels and the share of the population reaching basic proficiency, rather than on dispersion or inequality of skills (*Hanushek & Woessmann, 2008*).

A second line of research connects growth with human capital through technology adoption and diffusion. *Asif and Lahiri* treat human capital as multidimensional (distinguishing cognitive skills from years of schooling, life expectancy, and a learning-by-doing channel captured by lagged technology use) and find that the technology-human-capital link is conditional on both technology type and human-capital type. This is important because a nation's capacity to adopt and utilize new technologies may contribute to the economic value of skills in addition to their direct impact on output (*Asif & Lahiri, 2021*).

The relationship between human capital, growth, and inequality is further complicated by the fact that the distribution of qualifications shapes the distribution of market incomes. *Andersen* argues that human capital formation is how inequality and growth are connected, and that social barriers and capital-market inefficiencies can make inequality itself a barrier to investment in education. Unequal access to human capital formation results in a larger spread of qualifications, leading to a wider distribution of market income and weaker economic prospects. Using a stylized framework and cross-country evidence, *Andersen* highlights that countries with more equal educational distributions tend to achieve more equal disposable-income outcomes. Andersen (2015) argues that education policy, because it shapes both the level and distribution of qualifications, should be treated as a first-order redistributive tool, not merely a supplement to tax-and-transfer systems, and that raising the quality of schooling for the broad workforce matters at least as much as educating elites. (*Andersen, 2015*).

2.3. Technology diffusion and development

The technology-diffusion literature distinguishes clearly between invention at the frontier and the spread of existing technologies across countries. *Asif and Lahiri* argue that this distinction is central for growth analysis because traditional proxies such as total factor productivity do not directly measure adoption. They highlight that diffusion has an extensive margin, referring to whether and when a technology is adopted, and an intensive margin, referring to how intensively it is used once adopted. They contend that empirical research on

development should include direct measures of technology diffusion rather than relying solely on residual productivity indicators (*Asif & Lahiri, 2021*). Their analysis also demonstrates that diffusion is not consistent across nations or technologies. *Asif and Lahiri* model technology adoption using usage intensity and usage latency with the CHAT dataset, treating human capital as multidimensional. A main finding is that the relationship between technological spread and human capital is conditional, depending on both the type of technology and the type of human capital. Learning-by-doing is the most significant factor influencing diffusion, but cognitive abilities are more impactful than years of education or life expectancy in numerous models. The learning-by-doing effect is more pronounced in emerging countries, logical for follower economies that depend more on absorbing existing technologies than on pioneering frontier innovation. (*Asif & Lahiri, 2021*)

The World Bank distinguishes between AI invention, AI adaptation, and AI adoption. Innovation refers to the creation of new models, methods, and core applications; adaptation to modifying AI tools to local contexts; and adoption to the practical use of existing AI systems by firms, governments, or individuals. Most developing countries may benefit from AI without creating their own frontier models, gaining more from localized adaptation and the use of existing open-source platforms, especially where local languages, institutions, and sectoral needs differ from frontier economies. (*World Bank, 2025*)

However, the implications of AI diffusion for development can differ from earlier technologies. *Korinek and Stiglitz* argue that AI and related automation technologies are likely to be labor-saving and resource-saving and may weaken the traditional comparative advantages of many developing countries, creating the possibility of winner-takes-all dynamics, deteriorating terms of trade for labor-abundant economies, and even a reversal of convergence toward higher-income countries. This result is not certain. Developing nations can still mold their development strategies by directing technology adoption toward their own needs, including labor-using applications and broader sectoral strategies beyond manufacturing, since the direction of technology adoption is endogenous (*Korinek & Stiglitz, 2021*).

2.4. Skill distribution, inequality and inclusive growth

The preceding sections establish that AI operates through task-level mechanisms, cognitive skills robustly determine long-run growth, and technology diffusion is conditional on the type of human capital available. A separate question is whether the distribution of skills shapes the relationship between technological change and inequality, and whether productivity gains translate into broadly shared improvements in living standards.

Galor and Zeira (1993) develop an overlapping-generations model in which individuals differ in inherited wealth but not in potential skills. With credit market imperfections and indivisibilities in human capital investment, wealth distribution determines how many can invest in human capital each generation. Dynasties polarise into a skilled group that invests and an unskilled group that does not, with relative sizes determined by initial wealth distribution. The model produces multiple steady states rather than convergence, implying that historically determined inequality generates persistent cross-country differences in per-capita output. Countries where cognitive skills concentrate in a narrow segment of the

population may be structurally less able to translate new technologies into broad-based income gains (*Galor & Zeira, 1993*).

Castelló-Climent and Doménech (2002) construct education Gini coefficients for 108 countries over 1960-2000 and find that human capital inequality is a more robust predictor of growth than income inequality. The negative relationship survives regional dummies and is associated with lower investment rates, providing a transmission channel beyond direct productivity effects. *De Gregorio and Lee (2002)*, using a panel of more than 100 countries from 1960 to 1990, confirm that both higher average attainment and more equal education distribution are associated with lower income inequality. Greater dispersion of schooling significantly increases income inequality. These results establish the educational composition of the population as a central determinant of distributional outcomes (*Castelló-Climent & Doménech, 2002; De Gregorio & Lee, 2002*).

Whether technology-driven growth is inclusive depends further on the match between workforce skills and the type of technological change. Using task-level US data from 1980 to 2016, *Acemoglu and Restrepo (2022)* show that 50-70 % of the rise in wage inequality is accounted for by the displacement of workers specialised in routine tasks in rapidly automating industries. These large distributional effects coexist with modest aggregate productivity gains. This implies that automation-style technological change can produce significant inequality without proportionate growth, and the consequences depend on which workers are displaced, which depends on the skill composition of the workforce (*Acemoglu & Restrepo, 2022*).

At the global level, the distributional risks of AI extend to between-country divergence. The ILO and UN Tech Envoy (2024) document an emerging AI divide in which high-income nations benefit disproportionately while low- and medium-income countries lag due to gaps in infrastructure, investment, and skills. Many workers in developing countries lack basic digital tools, women face disproportionate exposure to generative AI's automating effects due to over-representation in clerical occupations, and a virtual brain drain may draw skilled professionals toward global platforms rather than national AI industries. (*ILO & UN Tech Envoy, 2024*).

2.5. Research gap

The literatures reviewed above identify the relevant mechanisms individually but have not brought them together in a single empirical framework applied to AI. The task-based models treat AI's productivity parameters as fixed constants calibrated to the United States, without allowing them to vary with the skill composition of the workforce. The human capital and diffusion literatures establish that cognitive skills condition technology absorption but have not applied this insight to AI adoption specifically. The inequality literature documents that automation can widen income gaps even when aggregate productivity rises but does not test whether cognitive skills moderate this channel across countries.

No existing study directly examines whether the cross-country productivity effects of AI readiness depend on the cognitive skill level of the population, or whether any such interaction carries through to poverty and distributional outcomes. This thesis fills the gap by

extending Acemoglu's (2024) task-based framework with a country-specific cognitive skill parameter and testing the resulting predictions on a cross-section of 98 countries.

3. Theoretical framework

3.1. Baseline task-based AI Model

The theoretical framework of this thesis builds on Acemoglu's (2024) task-based model of artificial intelligence. The model situates AI within a broader tradition of task-based technological change in which output is produced through a continuum of tasks that can be allocated either to labor or to capital.

3.1.1. Production Structure

The economy produces a single final good by combining a set of tasks indexed by $z \in [0, N]$, where N denotes the total measure of tasks available in the economy. The production function takes the form:

$$Y = B(N) * \left[\int_0^N y(z)^{\frac{\sigma-1}{\sigma}} dz \right]^{\frac{\sigma}{\sigma-1}}$$

where $y(z)$ is the output of task z . $\sigma > 0$ is the elasticity of substitution between tasks. When $\sigma < 1$, tasks are gross complements. $B(N)$ captures potential system-wide effects of new tasks.

Following Acemoglu (2024) and consistent with empirical estimates by Humlum (2021), the elasticity of substitution is set to $\sigma \approx 0.5$, implying that tasks are gross complements. Each task can be performed by either capital or labor, according to the task-level production function:

$$y(z) = A_L \cdot aL(z) \cdot l(z) + A_K \cdot aK(z) \cdot k(z)$$

where A_L and A_K are labor- and capital-augmenting productivity terms, $aL(z)$ and $aK(z)$ are task-specific productivity schedules for labor and capital respectively, And $l(z)$ and $k(z)$ denote labor and capital allocated to task z . A key structural assumption is that $\frac{aL(z)}{aK(z)}$ is increasing in z . This implies the existence of a threshold I such that all tasks $z \leq I$ are performed by capital and all tasks $z > I$ are performed by labor. Automation, in this framework, corresponds to an increase in I .

3.1.2. Four Channels of AI's Economic Impact

Within this framework, Acemoglu (2024) identifies four distinct channels through which AI can affect production. The first operates through extensive-margin automation, whereby AI reduces the cost of capital or increases capital's task-specific productivity $aK(z)$ for tasks near the automation threshold I . The second channel concerns task complementarities, through which AI raises the productivity of labor in tasks it does not fully automate, captured formally by an increase in $aL(z)$ or in the labor-augmenting productivity term A_L . This channel enables workers to specialize in higher-value activities within the same

occupation. The third channel, deepening of automation within already-automated tasks, is regarded by *Acemoglu* as less central to current generative AI. The fourth channel is new task creation, corresponding to an expansion of the task index N through the emergence of entirely new labor-intensive activities. As *Acemoglu (2024)* demonstrates, new task creation raises wages and tends to compress the gap between capital and labor income.

3.1.3. Hulten's Theorem and Aggregate Productivity

The central analytical result of *Acemoglu's (2024)* framework is that, in a competitive economy with constant returns to scale, the aggregate productivity effects of AI are governed by a version of Hulten's theorem. Regardless of which specific channel AI operates through (automation or task complementarity) the change in total factor productivity can be expressed as:

$$d \ln TFP = \int_N^0 \chi(z) \cdot \pi_{L(z)} dz$$

Where $\chi(z) = \frac{p(z)y(z)}{Y}$ is the GDP share of task z ,

and $\pi_{L(z)} = d \ln y(z)$ denotes the productivity improvement or cost savings in task z attributable to AI.

The total effect on GDP additionally incorporates the capital accumulation response:

$$d \ln = d \ln TFP + s_K * d \ln K$$

Where s_K is the capital share of GDP. Aggregating across tasks and denoting the economy-wide average cost savings by $\bar{\pi}$, the theorem simplifies to:

$$d \ln TFP = \bar{\pi} * (\text{GDP share of tasks impacted by AI})$$

Aggregate productivity gains from AI depend on precisely two quantities: the fraction of economic activity affected by AI, and the magnitude of task-level cost savings. Both are empirically bounded and considerably more modest than many public forecasts suggest.

3.1.4. Quantitative Implications

Combining estimates of AI's task exposure (*Eloundou et al., 2023*), the share of tasks cost-effective to automate (*Svanberg et al., 2024*), and labour cost savings from generative AI (*Noy & Zhang, 2023; Brynjolfsson et al., 2023*), *Acemoglu (2024)* arrives at an implied ten-year TFP gain of approximately 0.71%, or roughly 0.07 percentage points annually. Even under more optimistic assumptions the upper bound remains below 1%, and accounting for task difficulty reduces it further to about 0.55%.

3.1.5. Limitations of the Baseline Model and Motivation for Extension

The model is calibrated to the United States economy, drawing on O*NET occupational task data and wage estimates from the US Bureau of Labor Statistics. Its two-type labor structure treats the workforce as internally homogeneous within each skill group, thereby abstracting from the full distribution of cognitive abilities both within and across

countries. The parameters $\bar{\pi}$ and χ (average cost savings and the GDP share of tasks affected by AI) enter the model as fixed, economy-wide constants.

If the capacity of workers to complement, adopt, and productively engage with AI tools is conditioned by their cognitive capabilities (as suggested by the *World Bank's (2025)* competency condition and by *Hanushek and Woessmann's (2008)* evidence on the relationship between cognitive skills and economic growth) then both $\bar{\pi}$ and χ should be expected to vary systematically across countries with different skill distributions. An economy in which a substantial share of the workforce lacks the foundational cognitive skills required to operate AI-augmented tasks effectively will, ceteris paribus, realize smaller task-level productivity gains and a narrower effective scope of AI impact, even under equivalent levels of technology adoption.

3.2. Extending the model with cognitive skill level

3.2.1. Cognitive skill level as a modifier of AI productivity

3.2.1.1. Introducing the cognitive skill parameter

The extension introduces a country-specific cognitive skill parameter $h_c \in [0,1]$, defined as the share of the working-age population in country c that meets or exceeds a basic cognitive proficiency threshold (=basic literacy and numeracy benchmark used by *Hanushek and Woessmann (2008)*). A higher h_c indicates a more cognitively capable workforce on average. Workers below this threshold are assumed unable to realize meaningful productivity gains from AI complementarity.

While the model motivates h_c as a threshold-based share, the empirical analysis uses the normalised average HLO score. The two are monotonically related cross-nationally (countries with higher average scores consistently have larger proficient shares) so the model's qualitative predictions carry through. The continuous score is chosen for data availability and statistical efficiency, and results are unlikely to be sensitive to this distinction.

3.2.1.2. The skill utilization function

The effective task-level productivity gain realised in country c is modelled as the frontier gain scaled by a skill utilisation function $\varphi(h_c) \in [0,1]$, where φ is increasing in h_c . A workforce with no proficient workers realises no gain ($\varphi(0) = 0$), a fully proficient workforce realises the full frontier gain ($\varphi(1) = 1$), and intermediate values are strictly increasing. The realised gain is therefore $\pi(z, h_c) = \pi_0(z) \cdot \varphi(h_c)$.

3.2.1.3. Reinterpreting $\bar{\pi}$ and χ as country-specific

Returning to *Acemoglu's (2024)* aggregate result from Section 3.1.3:

$$d \ln TFP = \bar{\pi} \cdot \chi$$

the extension treats both $\bar{\pi}$ and χ as country specific. The average effective task-level productivity gain in country c is:

$$\bar{\pi}_c = \bar{\pi}_0 \cdot \varphi(h_c)$$

Similarly, the effective GDP share of tasks impacted by AI depends on h_c . In a country where a large share of workers fall below the proficiency threshold, many tasks nominally exposed to AI cannot in practice be augmented. This gives:

$$\chi_c = \chi_0 \cdot \psi(h_c)$$

where $\psi(h_c)$ is an analogous increasing function. The aggregate TFP result for country c therefore becomes:

$$d \ln TFP_c = \bar{\pi}_c \cdot \chi_c = \bar{\pi}_0 \cdot \varphi(h_c) \cdot \chi_0 \cdot \psi(h_c)$$

Two otherwise identical countries will realise different aggregate productivity gains from AI if they differ in h_c through both higher task-level productivity gains and a broader effective scope of AI impact.

3.2.2. Formalizing the skill-adjusted productivity function

Building on the country-specific application of Hulten's theorem from Section 3.2.1, two assumptions are introduced to make the framework empirically operational.

Assumption 1: The share of tasks covered by AI in country c is proportional to its level of AI readiness: $\chi_c = R_c$, where $R_c \in [0,1]$ is measured by the Oxford Insights Government AI Readiness Index, normalised to $[0,1]$.

Assumption 2: The average productivity gain per automated task is increasing in the cognitive capability of the workforce: $\bar{\pi}_c = h_c^\gamma$, where $h_c \in [0,1]$ is the normalised HLO score and $\gamma > 0$ is the skill elasticity of AI productivity.

Substituting both assumptions into $d \ln TFP_c = \bar{\pi}_c \cdot \chi_c$ and defining Γ as a global scaling constant yields the core expression:

$$d \ln TFP_c = \Gamma \cdot R_c \cdot h_c^\gamma$$

This is the skill-adjusted productivity function. Γ captures the frontier productivity gain from AI calibrated to *Acemoglu's (2024)* estimate of approximately 0.71% over a ten-year horizon. R_c captures cross-country variation in AI adoption and h_c^γ captures the cognitive amplification mechanism.

3.2.3. Testable Implications

The productivity function $d \ln TFP_c = \Gamma \cdot R_c \cdot h_c^\gamma$ generates three empirical predictions, each directly mapped to a model estimated in Section 4. The cross-country regressions identify conditional associations rather than structural causal effects (see Section 6). Instrumental variable approaches are acknowledged as a limitation and beyond the scope of this thesis.

The three hypotheses are nested rather than independent. H1 is the primary test of the skill-complementarity mechanism. H2 and H3 are conditional extensions that probe how any complementarity manifests across the skill distribution and whether it carries through to distributional outcomes. Because H2 and H3 presuppose the interaction mechanism that H1

tests, they are logically downstream of H1. If the interaction coefficient in M2 is statistically indistinguishable from zero, the mechanism that motivates the threshold and distributional predictions is not active in the data, and H2 and H3 are expected to return nulls as well. They are nevertheless tested and reported for completeness.

H1: Skill Complementarity (Primary Hypothesis)

The marginal return to AI readiness:

$$\frac{\partial(d \ln TFP_c)}{\partial R_c} = \Gamma \cdot h_c^\gamma$$

is strictly increasing in h_c . This is operationalised as a positive and statistically significant coefficient on the interaction term $R_c \times h_c$ in the baseline TFP growth regression (M2). This is the primary hypothesis of the thesis.

H2: Threshold Non-Linearity (Exploratory)

Conditional on H1 being supported, the growth effect of AI readiness should be economically meaningful among countries with higher cognitive skills and negligible among countries with lower cognitive skills. The prediction is that the relationship between AI readiness and TFP growth is systematically steeper (more positive) at higher levels of cognitive skills than at lower levels.

H3: Distributional Channel (Auxiliary)

Conditional on H1 being supported, if AI adoption raises growth more strongly in high- h_c countries, and if growth translates into poverty reduction through distributional channels, the poverty-reducing and income-distributional effects of AI adoption should also be stronger in higher-skill countries. The prediction is that, in countries with higher cognitive skills, AI readiness is associated with lower poverty and a larger income share accruing to the bottom of the distribution; in countries with lower cognitive skills, this association is weaker or absent.

4. Data and empirical strategy

4.1. Data sources

The empirical analysis draws on five external databases, all of which supply the analytical variables used in the main models. None of the robustness checks introduced in Section 4.6 (R4, R5, R16) requires additional data: R4 and R5 re-estimate M2 on restricted subsamples of the same estimation sample, and R16 re-uses the TFP series from Penn World Table 11.0 over an alternative 2015–2019 window. Table 4.1 summarizes the primary data sources.

Source	Publisher	Indicators Used	Analytical Role
Government AI Readiness Index (GAIRI), 2022 edition	Oxford Insights	Total score	R_c (Section 4.2)
Human Capital Project: Harmonised Learning Outcomes (HCP/HLO)	Oxford Insights	HLO test score	h_c (Section 4.3)

Penn World Table 11.0	University of Groningen	rtfpna	TFP outcome (primary); initial TFP level (convergence control); R16 temporal placebo window (Section 4.6)
Worldwide Governance Indicators (WGI)	World Bank	Six governance dimensions	Control (wgi_{avg})
World Development Indicators (WDI)	World Bank	DP per capita, poverty headcount, bottom-20 % income share, trade, investment, inflation, population growth, education spending; analytical regional classification	Log initial-income control; M3/M3b distributional outcomes; macro controls; regional classification for region FE

Table 4.1. Data sources

The Government AI Readiness Index (GAIRI), published by Oxford Insights annually since 2017, scores countries on the capacity of their public sector to adopt and govern AI. The 2022 edition, used here for R_c , covers 181 countries (*Oxford Insights, 2022*).

The Human Capital Project's Harmonised Learning Outcomes database constructs a cross-country measure of student achievement by rescaling region-specific assessments (PISA, TIMSS, PIRLS, EGRA and regional programmes) onto a single PISA-equivalent scale. The HLO scores span 2000–2024 for up to 173 economies. The 2020 cross-section is used as the primary value of h_c , with last-observation-carried-forward imputation for any country missing that year (*World Bank, 2023*).

Penn World Table 11.0 provides the TFP series for the primary outcome. $rtfpna$ (TFP at constant national prices, available for most countries through 2023) supplies both the annualised growth rate that enters M1-M2 and the initial TFP level used as a convergence control. The same series is re-used over the 2015–2019 window as the outcome in the temporal placebo R16 (Section 4.6) (*Feenstra, R. C. et.al, 2024*).

The Worldwide Governance Indicators report six dimensions (Voice and Accountability, Political Stability, Government Effectiveness, Regulatory Quality, Rule of Law, and Control of Corruption) as bounded scores in the interval $[-2.5, +2.5]$. Their arithmetic mean (wgi_{avg}) enters every specification as the institutional-quality control (*Kaufmann, D., Kraay, A., & Mastruzzi, M., 2010*).

The World Development Indicators database supplies the log initial-income control (NY.GDP.PCAP.KD, 2019) and the distributional outcomes for M3 and M3b (SI.POV.NAHC and SI.DST.FRST.20 respectively). It also supplies five macro controls (trade openness, investment share, CPI inflation, population growth, and government education expenditure), as well as the World Bank analytical regional classification used to construct the region fixed effects in M2, M3, and M3b (*World Bank, 2024*).

4.2. Measurement of AI Readiness R_c

4.2.1. Analytical variable

The 2022 GAIRI total score is structured into three pillars, each composed of several dimensions (*Oxford Insights, 2024*). The Government pillar captures national AI strategy (vision), regulatory and ethical frameworks (governance and ethics), and public administration skills and adaptability (internal digital capacity). The Technology Sector pillar measures innovation capacity through R&D spending and the business environment, technology market maturity, and human capital in the technology workforce. The Data and Infrastructure pillar assesses open data and statistics availability, population representativeness, and physical compute and connectivity infrastructure. The 2022, 2023, and 2024 editions share this three-pillar, ten-dimension framework, with 39-40 indicators per edition. The 2025 edition restructured the index substantially into six pillars, 14 dimensions, and 69 indicators (*Oxford Insights, 2025*). Only the headline total score, published continuously across all editions, is therefore used and the 2025 edition is not used.

The mapping from the raw 2022 total score to the analytical variable R_c follows three operations in the code. First, the 2022 cross-section is loaded, the total score extracted, and country names harmonised to ISO3 code. Second, R_c is min-max rescaled over the estimation sample, that is, over the set of countries that survive the merge with PWT 11.0 and the macro controls, so that

$$R_c = \frac{(GAIRI_{c,2020} - \min(GAIRI))}{(\max(GAIRI) - \min(GAIRI))} \in [0,1]$$

satisfying the unit-interval restriction of Assumption 1 in 3.2.2. Third, for the interaction specification in M2, R_c is mean-centred as $R_c^c = R_c - \bar{R}_c$ and the interaction term is constructed as $R_c^c * h_c^c$, with $h_c^c = h_c - \bar{h}_c$ so that the main coefficients admit the standard interpretation at sample means.

4.3. Measurement of cognitive skill level h_c

4.3.1. Data source and primary measure

The cognitive skill variable h_c is constructed from the World Bank's Harmonised Learning Outcomes (HLO) database. The analytical value of h_c is taken from the 2020 cross-section of the HLO database, the most recent pre-pandemic wave with broad country coverage. The 2020 cutoff captures the skill base immediately before the productivity outcome window, which begins in 2022 with the public release of large language models. Since cognitive skills in the working-age population change slowly, measuring them in 2020 to predict outcomes from 2022 onwards poses no material timing problem.

4.3.2. Normalization and temporal Imputation

The analytical variable h_c is obtained by min-max normalising the raw HLO scores over the estimation sample (the countries that survive the merge with PWT 11.0 and the macro controls) so that

$$h_c = \frac{(HLO_{c,2020} - \min(HLO))}{(\max(HLO) - \min(HLO))} \in [0,1]$$

For the interaction specification in M2, h_c is mean-centred analogously to R_c (Section 4.2.1), so that the same sample-mean interpretation applies to both main effects.

A small number of countries lack a 2020 HLO observation. To avoid sample loss, the series is extended over 2000-2025 and last-observation-carried-forward (LOCF) imputation is applied within country. A country missing the 2020 wave inherits its most recent prior HLO score, where one exists. Imputation is strictly within country and forward in time, never bridging across countries or carrying scores backward. Countries with no HLO observation at any point in the panel are dropped.

4.4. Outcome variables and controls

4.4.1. Primary Outcome: annualised TFP Growth

The estimation window spans 2022-2023, beginning in 2022, when large language models became publicly available, and ending in 2023, the most recent year for which *rtfpna* is available. Annualisation uses the compound-growth formula:

$$tfp_{growth_{avg}} = \left(\left(\frac{rtfpna_{end}}{rtfpna_{start}} \right)^{\frac{1}{T}} - 1 \right) \times 100$$

where *rtfpna_{start}* and *rtfpna_{end}* denote the values of *rtfpna* at the start and end years of the window and *T* is the length of the window in years. The series is multiplied by 100, so the estimated coefficients are expressed in percentage points of annual TFP growth. With *T* = 1, the formula reduces to the percentage change between 2022 and 2023. Should the 2023 endpoint be unavailable for enough countries, the analysis falls back to a three-year 2019-2022 window with an analogous annualised specification. This fallback was not invoked for the present sample. The outcome is winsorised at the 2nd and 98th percentiles.

The level of *rtfpna* at the start of the outcome window, *tfp_{level_{start}}*, is preserved separately and enters the conditioning set as a convergence control, because TFP growth exhibits conditional mean reversion: high-productivity economies grow more slowly as the gap to the frontier narrows, while follower economies grow faster from a lower base.

4.4.2. Conditioning variables

The conditioning set contains eight variables: two convergence controls and six macroeconomic controls. The convergence controls absorb the mechanical tendency of less-advanced economies to grow faster. The six macro controls account for structural country characteristics that affect productivity growth independently of AI readiness.

The six macro controls are averaged over the years immediately preceding the outcome window, with the endpoint year set to 2021, so the conditioning information reflects structural conditions before the AI-diffusion window begins. Five of them (the WGI composite, trade openness, investment, population growth, and inflation) are averaged over 2019-2021. Government education expenditure is averaged over a longer 2015-2021 window, as the series is published less frequently, has larger country-coverage gaps, and is slow-moving. Averaging over multiple years also smooths year-to-year volatility in series such as inflation or investment. The two convergence controls, by contrast, are single-year snapshots:

log GDP per capita in 2019 and the initial TFP level $tfp_{level_{start}}$ at the start of the outcome window (Section 4.4.1).

The two convergence controls are log real GDP per capita in 2019 and the initial TFP level $tfp_{level_{start}}$. Log initial income accounts for convergence toward the steady-state capital stock, the initial TFP level for convergence toward the efficiency frontier. Both correlate with development level, as do R_c and h_c by construction.

The six macro controls are the Worldwide Governance Indicators composite (wgi_{avg}), trade openness, investment as a share of GDP, government education expenditure, population growth, and CPI inflation. WGI captures institutional quality. Secure property rights, regulatory quality, and government effectiveness raise the elasticity of output to technological adoption, and countries with stronger institutions also tend to score higher on the GAIIRI for reasons unrelated to the cognitive-skill mechanism, biasing the R_c coefficient upward if omitted. Trade openness captures the technology-diffusion channel. More open economies absorb imported capital goods and frontier practices faster, and because the institutional environment producing high AI-readiness scores also tends to produce high trade integration, controlling for trade separates the AI-readiness channel from this broader diffusion mechanism. Investment as a share of GDP controls for capital deepening. Although TFP is conceptually separate from factor accumulation, the two are empirically correlated through the output-growth identity. Government education expenditure controls for public investment in human capital, separately from the cognitive achievement level h_c . Because high AI-readiness scores may correlate with broad public education spending, omitting it would conflate the cognitive-complementarity mechanism with the policy environment that produced it. Population growth captures demographic dilution. Faster growth raises the rate at which the capital stock must expand to hold capital per worker constant, shaping both investment costs and workforce absorptive capacity. CPI inflation captures macroeconomic stability. High and variable inflation disrupts long-horizon investment, distorts relative prices, and lowers allocative efficiency. Investment and inflation are winsorised at the 1st and 99th percentiles of the merged sample.

4.5. Empirical strategy

The analysis uses a cross-sectional OLS specification, exploiting between-country variation in AI readiness, cognitive skill, and post-2022 TFP growth to test H1-H3. The cross-section is the baseline because the theoretical model of Section 3.2.2 is a long-run cross-country relationship.

4.5.1. Cross-sectional baseline

The theoretical relation of 3.2.2,

$$\Delta \ln TFP_c \approx \Gamma \cdot R_c \cdot h_c^\gamma$$

is implemented as a linear-interaction specification. The cross-section does not separately identify the skill elasticity γ , the interaction analogue captures the model's central prediction

that the marginal return to AI readiness is increasing in the cognitive skill of the workforce, without requiring its recovery.

The additive baseline M1 establishes the unconditional association of each predictor with TFP growth, without interaction:

$$\Delta \ln TFP_c = \beta_0 + \beta_R R_c + \beta_h h_c + \beta'_X X_c + \varepsilon_c$$

M1 tests whether AI readiness and cognitive skill are individually associated with TFP growth after conditioning on the macro controls. Because the theoretical model predicts a multiplicative relationship between R_c and h_c , M1 is expected to understate their explanatory power if complementarity is present.

The interaction model M2 is the primary test of H1. It adds to M1 the product of the two mean-centred predictors and World Bank region fixed effects:

$$\Delta \ln TFP_c = \beta_0 + \beta_R R_c^c + \beta_h h_c^c + \beta_{Rh}(R_c^c * h_c^c) + \beta'_X X_c + \mu_{\text{region}(c)} + \varepsilon_c$$

Where $R_c^c = R_c - \bar{R}_c$ and $h_c^c = h_c - \bar{h}_c$, so β_R and β_h are the conditional associations evaluated at the sample mean of the other variable. The region fixed effects absorb broad geographic variation, confining identification of β_{Rh} to within-region differences. A positive, significant β_{Rh} would imply that the marginal return to AI readiness is increasing in cognitive skill, supporting H1. M2 uses the same sample as M1, so any change in fit is attributable to the interaction term and region dummies rather than to a change in country composition.

M3, replaces the TFP growth outcome with the national poverty headcount ratio (SI.POV.NAHC, averaged over 2015-2021) and keeps the M2 right-hand side:

$$Poverty_c = \beta_0 + \beta_R R_c^c + \beta_h h_c^c + \beta_{Rh}(R_c^c * h_c^c) + \beta'_X X_c + \varepsilon_c$$

It tests H3, the distributional channel: a negative and significant β_{Rh} would indicate that cognitive skills amplify the poverty-reducing association of AI readiness. Because poverty data cover only a subset of the estimation sample, the effective N is smaller than in M1-M2.

M3b, uses the income share of the bottom 20 percent (SI.DST.FRST.20, averaged over 2015-2021) as the dependent variable, again retaining the M2 specification:

$$IncShare_{B20,c} = \beta_0 + \beta_R R_c^c + \beta_h h_c^c + \beta_{Rh}(R_c^c * h_c^c) + \beta'_X X_c + \varepsilon_c$$

A positive β_{Rh} would signal that AI readiness is associated with a larger bottom-quintile income share in countries with higher cognitive skills. Together, M3 and M3b probe whether the growth-level complementarity documented in M2 translates into distributional gains, or whether it operates through channels that do not reach the bottom of the income distribution.

4.5.2. Inference

All models are estimated by OLS with HC3 heteroskedasticity-robust standard errors (Mackinnon & White, 1985). Significance is assessed at the 10%, 5%, and 1% levels. A Ramsey RESET test is applied to M2 to check functional-form adequacy. Standardised beta coefficients are reported for M1 and M2.

4.5.3. Diagnostics

Variance inflation factors, a Breusch-Pagan test, a Shapiro-Wilk test on residuals, and Cook's distance are reported in Section 5.2. Cook's-distance flagged observations are excluded in robustness check R4.

4.5.4. Sample construction and attrition

The estimation sample is constructed by left-joining the conditioning variables, HLO scores, and TFP outcomes onto the 181 countries covered by the 2022 GAIRI, then restricting to complete cases on all analysis variables. Three attrition counts are reported: the number of GAIRI countries, the subset with non-missing TFP growth at both window endpoints, and the final estimation sample.

To test whether attrition is systematic, a probit model is estimated in which the dependent variable equals one if a country lacks TFP data and zero otherwise. The explanatory variables are the AI readiness score and log GDP per capita. The latter is included because poorer countries are more likely to be missing from PWT 11.0 due to weaker statistical infrastructure.

4.6. Robustness checks

Following table depicts all robustness checks:

Check	Category	Description	Modification relative to M2
R4	Sample restriction	Exclusion of influential observations	Countries with Cook's distance > 4/N removed
R5	Sample restriction	Exclusion of transition economies	Former Soviet and Eastern European economies removed
R16	Temporal Placebo	Pre-AI placebo (2015-2019)	Dependent variable changed to TFP growth 2015-2019; 2022 R_c retained

Table 4.2. Robustness checks

Check R4 uses Cook's distance with the conventional 4/N threshold to flag observations exerting disproportionate leverage on the estimates, and re-estimates M2 after their exclusion. Check R5 removes former Soviet and Eastern European transition economies. The transition from plan to market involved severe output collapses and a switch from Soviet-era to UN national accounting standards (Campos & Coricelli, 2002), which can compromise the investment and capital stock series underlying PWT's perpetual-inventory-based TFP estimates (Inklaar & Timmer, 2013). Finally, a temporal placebo test rules out the possibility that the estimated interaction captures a pre-existing structural pattern. Check R16 applies the full M2 specification to TFP growth computed over 2015-2019, retaining the 2022 AI

readiness score as the predictor. A null result in the pre-AI window strengthens the interpretation that the interaction is particular to the period in which AI technologies began diffusing at scale.

5. Results

5.1. Sample description and descriptive statistics

The sample-construction procedure described in Section 4.5.4 yields a final regression sample of 98 countries for M1 and M2, with smaller subsamples of 77 (M3) and 89 (M3b) for the distributional models. Excluded countries are concentrated in Sub-Saharan Africa, East Asia and the Pacific, MENA, and LAC, and the estimation sample is therefore tilted toward higher-income, more AI-ready economies than the world average.

Figure 5.1 plots the two core predictors against one another.




Figure 5.1.

The two variables are strongly positively correlated ($r = 0.82$; see also Figure 5.4). Countries with the highest cognitive skill levels, like Singapore, the United States, and Western European economies such as the Netherlands, Finland, and Denmark, also score near the top of the readiness index, while Sub-Saharan and South Asian economies cluster in the lower left. This overlap constrains the interaction test. Once the two main effects are conditioned on, the M2 product term has little residual variation left to exploit. The same correlation shapes the sample layout: no low- h_c country appears in the top AI-readiness quintile, leaving the top-right cell of the quintile plot in §5.4 empty.

Table 5.1 reports descriptive statistics for the 98 countries in the regression sample.

	Unique	Missing Pct.	Mean	SD	Min	Median	Max
TFP Growth (%)	97	0	-0.5	2.4	-7.1	-0.4	5.1
AI Readiness (r_c)	98	0	0.5	0.2	0.1	0.5	1.0
Cog. Skills (h_c)	98	0	0.5	0.2	0.0	0.5	1.0
GDP per capita (constant 2015 US\$)	98	0	9.1	1.4	6.0	9.2	11.6
Governance (WGI)	98	0	0.2	0.9	-1.5	-0.0	1.9
Trade (% GDP)	98	0	89.5	57.5	25.2	76.9	385.2
Investment (% GDP)	98	0	24.0	6.8	8.0	22.6	42.9
Education Spending	98	0	14.7	4.7	6.6	13.7	28.1
Pop. Growth (%)	98	0	0.9	1.1	-1.5	0.9	3.3
Inflation (CPI)	98	0	5.2	12.8	-0.4	2.5	112.8
Initial TFP level	98	0	1.0	0.0	0.7	1.0	1.1

Table 5.1.

Mean annualised TFP growth over 2022-2023 is -0.5 % (SD 2.4 pp, range -7.1 to +5.1 % after winsorisation at the 2nd and 98th percentiles), the negative mean reflecting widespread productivity deceleration following the post-pandemic recovery. Both R_c and h_c span nearly the full [0, 1] normalised range (means ≈ 0.50 , SD ≈ 0.20), and log GDP per capita spans 6.0 to 11.6 (median 9.2), confirming that the sample includes both low- and high-income economies. Among the conditioning variables, two exhibit notable right skew: trade openness (median 76.9 %, max 385.2 %, capturing small open economies such as Singapore and Luxembourg) and CPI inflation (median 2.5 %, max 112.8 %, driven by a few crisis-economy observations). Governance (WGI) ranges from -1.5 to 1.9 (mean 0.2), and population growth concentrates in a narrow positive range (mean 0.9 %, SD 1.1). The remaining conditioning variables display the expected cross-country dispersion.

The distribution of the outcome variable and a regional decomposition illustrating the between-region variation absorbed by the World Bank region fixed effects in M2 are reported in Appendix A.1 (Figures A.1 and A.2).

5.2. Diagnostics

Three diagnostic checks are reported for the M2 specification: variance inflation factors to test for multicollinearity, a Ramsey RESET test for functional-form adequacy, and Breusch-Pagan tests for heteroscedasticity across M1, M2, and M3.

The highest VIF values are recorded for log initial GDP per capita (6.47) and governance (6.29). Centred AI readiness registers a VIF of 5.97 and centred cognitive skills of 4.27. The VIF for the interaction term $r_c \times h_c$ is only 1.32. The remaining controls fall comfortably below 2: trade (1.37), education spending (1.37), population growth (2.00), inflation (1.20), investment (1.15), and the initial TFP level (1.15).

Figure 5.4. shows the Pearson correlation matrix of all analysis variables.




Figure 5.4.

The strongest pairwise relationship is between AI readiness and log GDP per capita ($r = 0.87$), followed closely by log GDP per capita and governance ($r = 0.86$) and by AI readiness and governance ($r = 0.84$). AI readiness and cognitive skills themselves correlate at $r = 0.82$. Negative associations with population growth (-0.56 , -0.57) and education spending (-0.38 , -0.41) are also visible.

The test yields $F = 0.849$ ($p = 0.432$), failing to reject the null of correct specification at any conventional level.

The Breusch-Pagan test provides no evidence of heteroscedasticity in the TFP regressions: 8.084 ($p = 0.621$) for M1 and 14.128 ($p = 0.658$) for M2. For M3, the poverty headcount specification yields 30.861 ($p = 0.014$), rejecting homoscedasticity at the 5% level.

5.3. Baseline results M1 and M2 (H1 Test)

Table 5.2 reports the baseline OLS estimates for all four specifications. Columns (1) and (2) test H1 directly using the M1 and M2 specifications defined in §4.5.1. Columns (3) and (4) extend the same right-hand side to the distributional outcomes discussed in Section 5.5.

Regression Results: v6 (HC3 SEs; region dummies hidden)

Dependent variable: TFP Growth 2022-2023 (annualised %)

	M1	M2 + FE	M3 + FE	M3b + FE
	(1)	(2)	(3)	(4)
AI Readiness Index 2022 (rc)	-0.233 (2.231)			
AI Readiness Index (rc, centred)		-2.706 (3.244)	-11.111 (18.628)	-2.254 (1.834)
Cognitive Skills Index (hc)	4.394 (2.779)			
Cognitive Skills Index (hc, centred)		2.793 (2.743)	15.183 (14.233)	1.600 (1.584)
rc × hc		-2.327 (5.786)	95.432** (36.097)	4.570 (4.377)
Log GDP p.c. (2019)	-0.789 (0.564)	-0.836 (0.673)	-4.919* (2.909)	-0.645** (0.311)
Governance (WGI)	-0.634 (0.673)	-0.338 (0.772)	-4.421 (3.989)	1.070** (0.452)
Trade (% GDP)	-0.002 (0.005)	-0.002 (0.006)	0.039 (0.027)	-0.001 (0.003)
Investment (% GDP)	-0.005 (0.053)	-0.013 (0.060)	-0.143 (0.208)	0.013 (0.023)
Education Spending (%)	0.061 (0.078)	0.085 (0.084)	0.081 (0.333)	0.010 (0.039)
Population Growth (%)	-0.448 (0.350)	-0.143 (0.433)	-0.408 (1.823)	0.132 (0.203)
Inflation (CPI)	-0.033 (0.020)	-0.027 (0.022)	-0.047 (0.092)	0.005 (0.012)
Initial TFP level (2022)	1.108 (14.024)	0.172 (13.330)	-9.611 (31.231)	-7.098* (3.931)
Intercept	3.555 (15.287)	7.660 (15.657)	66.854* (36.390)	19.426** (4.364)
Observations	98	98	77	89
R ²	0.192	0.267	0.684	0.580
Adjusted R ²	0.100	0.111	0.600	0.479
World Bank region FE?	No	Yes	Yes	Yes

Notes: *p<0.1; **p<0.05; ***p<0.01. Robust (HC3) SEs in parentheses below coefficients. rc = 2022 Oxford AI Readiness Index (lagged). M1 includes raw rc and raw hc but not the interaction (no FE specification). M2 + FE, M3 + FE and M3b + FE include region fixed effects, with rc and hc reported in group-centred form (rc, centred / hc, centred) so that the main effects can be identified alongside the FE; the interaction rc × hc is reported for these three specifications only. Region dummies are omitted from display.

Table 5.2.

The additive specification M1 already foreshadows the central finding. In a sample of 98 countries neither AI readiness nor cognitive skills moves TFP growth in a statistically detectable way: the AI-readiness coefficient is essentially zero (-0.23, SE = 2.23) and the cognitive-skills coefficient is positive but imprecisely estimated (4.39, SE = 2.78). The model

explains roughly 19 per cent of the cross-country variation in TFP growth ($R^2 = 0.19$, adjusted $R^2 = 0.10$), but the substantive story sits in what does *not* emerge, that even conditioning variables that macro-growth theory would expect to matter, such as log initial GDP per capita and governance, fall short of conventional thresholds, and inflation only nudges the 10% line.

A sample this small, observed over a window this short, simply does not generate the precision needed to detect moderate associations. The null on the AI variables should therefore be read against that backdrop rather than as evidence of absence.

M2 sharpens the test by introducing the interaction and the region fixed effects. The interaction coefficient is -2.33 with a standard error of 5.79 (economically small and statistically indistinguishable from zero) while the centred main effects remain individually insignificant (-2.71 for AI readiness, 2.79 for cognitive skills). The R^2 rises from 0.19 to 0.27, but the gain is attributable almost entirely to the region dummies: the adjusted R^2 moves only from 0.10 to 0.11, so the interaction itself contributes no incremental explanatory power. This is consistent with what the diagnostics in Section 5.2 already foreshadowed. R_c and h_c co-move so tightly that once both main effects are conditioned on, the interaction has little independent information left to carry. The wide confidence bands around all three coefficients (Appendix A.2, Figure A.3) reinforce this: every interval comfortably spans zero, and the M2 standard errors widen proportionally to the centred main effects, leaving every t-ratio below one. Once the region fixed effects absorb the broad geographic gradient, what remains is a within-region slice of variation too narrow for a complementarity to register, even if one exists in the underlying data.

The baseline results therefore do not support H1. The data provide no evidence that AI readiness and cognitive skills interact to predict TFP growth over 2022-2023, and the null carries two substantive readings that the present design cannot separate. The first is genuinely early: generative AI tools only began diffusing at scale in late 2022, so by the end of the 2022-2023 growth window most economies had barely begun integrating them into production, and the cross-country variation in readiness a strategy-and-infrastructure proxy has yet to translate into variation in deployment. The second is mechanical power: a 98-country cross-section estimated over a single one-year growth window is dominated by year-to-year TFP noise, and the interaction absorbs further degrees of freedom on top of that.

5.4. Marginal effects and threshold probing (H2)

H2 asks whether the average null in M2 hides heterogeneity at specific skill levels. Two probes, reported in Appendix A.4, test this and converge on the same answer. The predicted marginal effect of AI readiness is flat whether evaluated at the 10th, 50th, or 90th percentile of cognitive skill, with confidence bands overlapping throughout (Figure A.5). A non-parametric quintile plot tells the same story. Neither group shows a monotonic AI-readiness-TFP-growth relationship, the two lines cross, and confidence intervals overlap at every quintile (Figure A.6). A bivariate scatter split at the h_c median (Figure A.4) corroborates this. The binned means are unstable, the cell sizes range from two to nineteen because no low- h_c country falls in the top AI-readiness quintile, a direct consequence of the strong correlation between h_c and $-R_c$.

AI readiness shows no significant association with TFP growth at any skill level. The evidence does not support H2.

5.5. Distributional Channel H3 (M3, M3b)

The distributional models test whether AI readiness translates into poverty reduction, and whether cognitive skills amplify or reshape that association. M3 (poverty headcount) produces the one significant AI-related coefficient in the entire study, but its sign runs opposite to what a complementarity reading of H3 would predict.

The interaction term in column (3) of Table 5.2 is positive and significant (95.43, SE = 36.10, $p < 0.05$), meaning the marginal association between AI readiness and poverty becomes *less* negative as cognitive skills rise. At the sample mean of h_c the association is indistinguishable from zero; one standard deviation below the mean, AI readiness is associated with roughly 30 percentage points lower poverty, while one standard deviation above the mean the effect reverses to a weakly positive value. AI readiness is thus associated with lower poverty only in countries with below-average cognitive skills, and the association weakens as human capital rises. Cognitive skills do not amplify the poverty-reducing association of AI readiness, but they attenuate it.

Figure 5.9 plots poverty headcount ratios against AI readiness, with separate linear fits for the high- h_c and low- h_c groups.




Figure 5.9

Both lines slope downward, but the low- h_c group shows a visibly steeper negative gradient, exactly as the positive interaction predicts. The high- h_c observations cluster in the lower-right quadrant (low poverty, moderate-to-high AI readiness), leaving the upper-right region of high-poverty, high-AI-readiness countries largely unoccupied, which is a feature of the development-cluster correlation rather than a causal effect of AI. Several African economies (CAF, MOZ, SLE, ZMB) anchor the high-poverty end of the low- h_c line, while European countries sit at the opposite corner.

M3b replaces the dependent variable with the income share of the bottom quintile. Here the AI-related coefficients are uniformly insignificant: the interaction is positive (4.57, SE = 4.38), directionally consistent with M3, but its magnitude relative to its standard error is too small to reject zero. Governance is the only robust predictor of the bottom-quintile share (1.07, SE = 0.45, $p < 0.05$), with log initial GDP per capita significantly negative (-0.65, SE = 0.31, $p < 0.05$). The bivariate scatter (Appendix A.5, Figure A.7) shows no statistically distinguishable gradient by h_c group, consistent with the insignificant M3b interaction. The contrast between M3 and M3b is itself informative. In absolute poverty, AI readiness carries a significant association, but conditional on human capital being low, not high. In relative inequality, no association survives. The distributional channel that H3 posits does not operate in the data. Where AI readiness matters for the poor, it does so despite low cognitive skills rather than because of high ones.

5.6. Robustness (R4, R5, R16)

Three sensitivity checks probe whether the M2 estimates are driven by groups of countries or by pre-existing structural patterns. Table A.2 reports the Cook's-distance exclusion (R4), Table A.3 the transition-economy exclusion (R5), and Table A.4 the temporal placebo on 2015-2019 TFP growth (R16), all reported in Appendix A.6. All three specifications retain the M2 right-hand side, including the centred interaction and World Bank region fixed effects.

Across all three robustness checks the AI-readiness $\times$ cognitive-skills interaction remains indistinguishable from zero, consistent with the precision limits discussed in §5.3. R4 attenuates its magnitude without troubling the standard errors, R5 attenuates it further, and R16 flips its sign. Cognitive skills, in contrast, emerges as the more systematic correlate. It is positive in every specification and crosses conventional thresholds when structurally noisy economies are excluded (R5) and remains positive but insignificant in the temporal placebo (R16).

5.7. Within-country identification check (Panel FE)

A within-country fixed-effects specification using annual variation in the GAIRI over 2022-2024 (Appendix Table A.5 and Figures A.8-A.9) asks whether the cross-sectional null in M2 survives once time-invariant country characteristics are partialled out. The answer is barely, and not robustly. Under country-only fixed effects the AI-readiness coefficient is negative and significant (-1.72, SE = 0.71, $p < 0.05$), but it loses significance once year fixed effects are added (-1.36, SE = 0.88) and weakens further when time-varying governance is included (-1.19, SE = 0.82). The demeaned bivariate scatter is essentially flat, and the within-country

standard deviation of AI readiness is very small (at most two to three annual observations per country) so much of what the panel identifies reflects transient noise rather than a structural within-country relationship. The cross-sectional null in M2 is therefore not driven solely by omitted country characteristics, but the panel cannot credibly separate any AI effect from common post-pandemic recovery dynamics over a window this short.

6. Limitations

The estimates are cross-sectional OLS with World Bank region fixed effects, so they identify from between-country variation at one point in time. Countries with higher AI readiness also differ in unobserved ways (sectoral composition, regulatory regime, legal origin, colonial history) that are not controlled for, and any of these could drive the coefficients through channels unrelated to the skill-AI complementarity the theory posits. The estimates should be read as conditional associations, not as causal effects.

The wide null established in §5.3 also constrains the panel check in §5.7, which adds no credible within-country identification over a two-to-three-year window.

Two measurement gaps further weaken the link between the empirical proxies and the theory. The 2022 Government AI Readiness Index scores national strategies, regulatory frameworks, and digital infrastructure, not actual AI use in production, so a country can score high without having integrated AI into firms' processes over 2022-2023. The cognitive-skills variable is constructed from HLO scores taken by school-age children, while the theory concerns the working-age population exposed to AI in production. The two are plausibly correlated across countries but the mapping is not verified. The model further defines h_c as the share above a proficiency threshold, while the empirical implementation uses the continuous, mean-normalised HLO score. The two are monotonically but not linearly related, and the lowest observed HLO is well above the theoretical zero, so point estimates at very low h_c are uncertain.

The distributional models add a temporal mismatch. M3 and M3b use poverty and bottom-quintile income shares averaged over 2015-2021, not changes over the 2022-2023 outcome window, so they speak to the association between AI readiness and the existing cross-country distribution of poverty rather than to whether AI adoption changed it.

Finally, the sample is not representative. Of the 181 countries in the GAIRI universe, only 98 enter M1/M2, the rest are lost at the PWT merge stage, mostly because Penn World Table coverage is thinner in poorer regions. In-sample mean AI readiness is 0.51 versus 0.27 for the excluded group, and mean log GDP per capita is 8.98 versus 8.09. The analysis is carried out on a set of countries that are richer and more AI-ready than the world average, while the poorest and least AI-ready countries are largely missing.

7. Discussion and Conclusion

This thesis asked whether the effect of AI readiness on total factor productivity growth depends on the cognitive skill level of the workforce. The cross-sectional evidence from 98

countries over 2022-2023 does not confirm the prediction. The interaction coefficient in M2 is indistinguishable from zero and remains so across all three robustness checks.

What does emerge more reliably is the direct association between cognitive skills and productivity growth. This pattern aligns with *Hanushek and Woessmann's (2008)* finding that the quality of human capital carries the durable cross-country relationship with growth, and with *Asif and Lahiri's (2021)* emphasis on cognitive abilities as a more systematic correlate of technology diffusion than years of schooling. In these results skills operate as a direct correlate of productivity growth rather than as the moderator of AI adoption that the extended model anticipated.

Because H2 and H3 are conditional on H1 and the H1 interaction is indistinguishable from zero, the downstream hypotheses are expected to return nulls as well, and they are reported for completeness. The threshold probing for H2 indeed converges on the same conclusion as the baseline: the parametric and non-parametric probes in Appendix A.4 show no systematic divergence between high- and low- h_c groups, and the small, uneven cell sizes leave the bin means too unstable to support a threshold reading. The model predicts heterogeneity by skill level, but the data do not display it.

The H3 result is the least expected. The significant interaction carries the opposite sign to complementarity, pointing to a poverty-floor artefact rather than a distributional channel. Because high-skill countries already sit at minimal poverty and all three variables share the same development cluster, the M3 coefficient is best read as a conditional cross-sectional association rather than evidence that AI readiness and cognitive skills jointly shape the income distribution.

The within-country panel check does not contradict the cross-sectional null. Its only significant coefficient is fragile, dissolving once year fixed effects are added, and with at most three annual observations per country the panel cannot separate AI effects from post-pandemic recovery dynamics.

These results do not rule out skill-AI complementarity, but they indicate that, at the current diffusion stage and with available measurement, the mechanism cannot be distinguished from zero at the cross-country level. Sharpening the test would require longer outcome windows as subsequent PWT releases permit, direct measures of AI deployment rather than readiness scores, adult skill assessments such as PIAAC rather than school-age HLO scores, and a sector or firm unit of analysis that exploits wider variation in AI exposure and skill intensity.

Until such data arrive, the finding that stands is the one the human capital literature would predict, that cognitive skills are associated with productivity growth on their own terms. The thesis therefore contributes a clean negative on the cross-country interaction hypothesis, a positive on the direct skill-growth relationship that survives perturbations to the sample and outcome window, and a caveat that the distributional associations for poverty reflect the development cluster rather than an AI-driven channel.

8. References

World Bank. (2025). Digital progress and trends report 2025: Strengthening AI foundations. World Bank. <https://doi.org/10.1596/978-1-4648-2264-3>

Microsoft AI Economy Institute. (2026, January). Global AI adoption in 2025: A widening digital divide. Microsoft.

Comunale, M., & Manera, A. (2024, March). The economic impacts and the regulation of AI: A review of the academic literature and policy actions (IMF Working Paper No. 24/65). International Monetary Fund.

OECD. (2025). Governing with artificial intelligence: The state of play and way forward in core government functions. OECD Publishing.

Hanushek, E. A., & Woessmann, L. (2008). The role of cognitive skills in economic development. *Journal of Economic Literature*, 46(3), 607–668.
<https://doi.org/10.1257/jel.46.3.607>

Acemoglu, D. (2024). *The simple macroeconomics of AI* NBER Working Paper 32487 (2024), <https://doi.org/10.3386/w32487>.

Aghion, P., Jones, B. F., & Jones, C. I. (2017). Artificial intelligence and economic growth (NBER Working Paper No. 23928). National Bureau of Economic Research.

Saam, M. (2024). Macroeconomic productivity effects of artificial intelligence. *The Economists' Voice*, 21(2), 327–333. <https://doi.org/10.1515/ev-2024-0072>

Asif, Z., & Lahiri, R. (2021). Dimensions of human capital and technological diffusion. *Empirical Economics*, 60, 941–967. <https://doi.org/10.1007/s00181-019-01777-3>

Andersen, T. M. (2015). Human capital, inequality and growth (European Economy Discussion Paper No. 007). European Commission. <https://doi.org/10.2765/7521>

Korinek, A., & Stiglitz, J. E. (2021). Artificial intelligence, globalization, and strategies for economic development (NBER Working Paper No. 28453). National Bureau of Economic Research. <https://doi.org/10.3386/w28453>

Campos, N. F., & Coricelli, F. (2002). Growth in transition: What we know, what we don't, and what we should. *Journal of Economic Literature*, 40(3), 793–836.
<https://doi.org/10.1257/002205102760273797>

Inklaar, R., & Timmer, M. P. (2013). Capital, labor and TFP in PWT 8.0 [Working paper]. University of Groningen, Groningen Growth and Development Centre.
https://www.rug.nl/ggdc/docs/capital_labor_and_tfp_in_pwt80.pdf

- Acemoglu, D., & Restrepo, P. (2022). Tasks, automation, and the rise in U.S. wage inequality. *Econometrica*, 90(5), 1973–2016. <https://doi.org/10.3982/ECTA19815>
- Castelló-Climent, A., & Doménech, R. (2002). Human capital inequality and economic growth: Some new evidence. *Economic Journal*, 112(478), 187–200.
- De Gregorio, J., & Lee, J.-W. (2002). Education and income inequality: New evidence from cross-country data. *Review of Income and Wealth*, 48(3), 395–416.
- Galor, O., & Zeira, J. (1993). Income distribution and macroeconomics. *Review of Economic Studies*, 60(1), 35–52.
- International Labour Organization & United Nations Office of the Secretary-General's Envoy on Technology. (2024). Mind the AI divide: Shaping a global perspective on the future of work. International Labour Office.
- Humlum, A. (2021). Robot adoption and labor market dynamics (Unpublished job market paper). Princeton University.
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. arXiv. <https://arxiv.org/abs/2303.10130>
- Svanberg, M., Akselrod, R., Li, V., Ford, B., & Brynjolfsson, E. (2024). Beyond AI exposure: Which tasks are cost-effective to automate with computer vision? (CES Working Paper No. 24-21). U.S. Census Bureau.
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381 187-192(2023). DOI: 10.1126/science.adh2586
- Brynjolfsson, E., Li, D., & Raymond, L. (2023). Generative AI at work NBER Working Paper 31161 (2023), <https://doi.org/10.3386/w31161>.
- MacKinnon, J. G., & White, H. (1985). Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *Journal of Econometrics*, 29(3), 305–325.
- Oxford Insights. (2022). Government AI readiness index 2022. Oxford Insights.
- Oxford Insights. (2023). Government AI readiness index 2023. Oxford Insights.
- Oxford Insights. (2024). Government AI readiness index 2024. Oxford Insights.
- Oxford Insights. (2025). Government AI readiness index 2025. Oxford Insights.
- World Bank. (2023). Human Capital Index and Harmonised Learning Outcomes.

World Bank Group, Human Capital Project.

<https://www.worldbank.org/en/publication/human-capital>

Feenstra, R. C., Inklaar, R., & Timmer, M. P. (2024). Penn World Table 11.0. University of Groningen, Groningen Growth and Development Centre.

<https://www.rug.nl/ggdc/productivity/pwt/>

Kaufmann, D., Kraay, A., & Mastruzzi, M. (2010). The Worldwide Governance Indicators: Methodology and Analytical Issues. World Bank Policy Research Working Paper No. 5430.

World Bank. (2024). World Development Indicators.

<https://databank.worldbank.org/source/world-development-indicators>

Appendix

This appendix collects the supplementary tables and figures referenced in Section 5: alternative presentations of the baseline and robustness results, the threshold-probing diagnostics for H2, the distributional scatter for M3b, and the within-country panel fixed-effects specifications.

A.1 Descriptive Supplements to §5.1.

Figure A.1 shows the distribution of the outcome variable across the 98-country estimation sample.




Figure A.1.

The histogram is concentrated within roughly ± 2 percentage points of zero, with the modal bar sitting just above the 0 % line and a longer left tail extending to the winsorised lower bound of -7.1 per cent.

Figure A.2 displays TFP growth by World Bank region, illustrating the between-region variation that is absorbed by the region fixed effects in M2.




Figure A.2.

Median growth hovers near or slightly below zero in every region; Latin America and the Caribbean is the single region whose median rests on the positive side, while the Middle East, North Africa, Afghanistan and Pakistan region and Sub-Saharan Africa record the lowest medians ($\approx -1\%$). The inter-quartile ranges are wide and overlap substantially across regions, and within Europe and Central Asia the whiskers and outliers together span roughly -7 to +5 percentage points, indicating that within-region dispersion dominates any between-region shift.

A.2. Standardised beta coefficients (M1, M2) - supplement to §5.3

The standardised beta coefficients (Table A.1) confirm that AI readiness and cognitive skills are not the dominant predictors of TFP growth in the sample.

Standardised Beta Coefficients (M1, M2 + FE) - v6

Dependent variable: tfp_growth_avg

	M1 Std. (1)	M2 Std. + FE (2)
rc (raw)	-0.023 (0.217)	
rc_c (centred)		-0.263 (0.316)
hc (raw)	0.398	

	(0.252)	
hc_c (centred)		0.253
		(0.249)
rc × hc (rc_hc)		-0.045
		(0.113)
log_gdp_pc_2019	-0.456	-0.484
	(0.326)	(0.389)
wgi_avg	-0.227	-0.121
	(0.241)	(0.276)
trade_pct_gdp	-0.045	-0.055
	(0.122)	(0.134)
inv_pct_gdp	-0.014	-0.035
	(0.149)	(0.168)
edu_spend_pct	0.117	0.163
	(0.149)	(0.162)
pop_growth_pct	-0.210	-0.067
	(0.164)	(0.203)
inflation_cpi	-0.173	-0.141
	(0.108)	(0.116)
tfp_level_start	0.018	0.003
	(0.223)	(0.212)
Constant	-0.000	0.442
	(0.107)	(0.312)
World Bank region	No	Yes
FE?		
Observations	98	98
R ²	0.192	0.267
Adjusted R ²	0.100	0.111

Notes: *p<0.1; **p<0.05; ***p<0.01. Numeric variables z-standardised (mean 0, SD 1); region kept as factor. Robust (HC3) SEs in parentheses below standardised coefficients. M1 includes raw rc and raw hc (no interaction); M2 + FE uses group-centred rc_c / hc_c and the centred interaction rc × hc, with region fixed effects.

Table A.1

In M1, a one-standard-deviation increase in h_c is associated with a 0.398-standard-deviation increase in TFP growth, the largest standardised effect among the key variables but statistically insignificant (SE = 0.252). The standardised coefficient on R_c is -0.023 (SE = 0.217), close to zero. In M2, the standardised interaction coefficient is -0.045 (SE = 0.113), the centred main effects move to -0.263 for R_c and 0.253 for h_c , both insignificant. Log initial GDP per capita carries the largest standardised coefficient among the controls in both specifications (-0.456 in M1, -0.484 in M2), but neither estimate achieves statistical significance at conventional levels. Inflation is the closest control to a conventional threshold but remains insignificant in both specifications (-0.173, SE = 0.108 in M1, -0.141 [SE = 0.116] in M2). No variable in either column attains significance, meaning that within this sample and time window, no single predictor systematically explains cross-country TFP growth differences with statistical confidence.

Figure A.3 visualises the coefficient estimates and 95 per cent confidence intervals for the key variables in M1 and M2 as a forest plot.




Figure A.3.

The wide confidence bands around R_c , h_c , and the interaction term illustrate the imprecision of the estimates, with all intervals comfortably spanning zero. The point estimates for the centred main effects in M2 are somewhat larger in absolute value than their M1 counterparts, but the standard errors widen proportionally, so the t-ratios remain below one in every case.

A.3 Binned scatter by h_c median - supplement to §5.4

Figure A.4 presents the bivariate scatter of AI readiness against TFP growth, with separate OLS regression lines for countries above and below the median cognitive skill level.





Figure A.4

The two regression lines are nearly flat and broadly parallel, with no visible divergence between the high- h_c and low- h_c slopes. Countries such as the United States, India, and Kenya cluster near or above the zero-growth line regardless of their AI readiness level, while Argentina stands out as a large negative outlier at the low- h_c end and Singapore as a mild negative outlier among high- h_c economies. The absence of a diverging pattern between the two groups is consistent with the insignificant interaction in M2: if cognitive skills enabled AI-driven productivity gains, one would expect the high- h_c line to slope upward more steeply than the low- h_c line, and no such pattern is visible.

A.4 Marginal-effects and quintile probes for H2 - supplement to §5.4

Figure A.5 plots predicted TFP growth as a function of mean-centred AI readiness at three representative levels of cognitive skills (10th, 50th, and 90th percentile), holding all other controls at their sample means and the region fixed effect at the modal region (Europe &

Central Asia).



∴ Region FE set to modal region (Europe & Central Asia). Controls at sample means. Shaded = 95% CI. rc = 2022 Oxford AI Readiness (lagged). TFP from PWT 11.0 (2022–2023).

Figure A.5

The three prediction lines slope gently downward across the observed range of R_c , and the vertical separation between them is modest: the low- h_c line sits slightly above the median and high- h_c lines at the left of the plot, but the ordering reverses toward the right without any consistent divergence. The 95 per cent confidence bands are wide and overlap substantially at every value of R_c , so the three predicted trajectories are statistically indistinguishable from each other and from a flat line at zero. The absence of any visible divergence among the three lines confirms that, across the observed range of AI readiness, predicted TFP growth does not differ meaningfully between low-, median-, and high-skill countries.

Figure A.6 shows a non-parametric perspective by computing mean TFP growth within each AI-readiness quintile, separately for countries above and below the median cognitive-skills level.




Figure A.6

If skill complementarity were present, the high- h_c group should exhibit a steeper positive slope across the quintiles than the low- h_c group. The observed pattern is inconsistent with this prediction. Neither group shows a monotonic relationship between AI readiness and TFP growth: the high- h_c line rises from -1.0 in the first quintile to 2.3 in the second, then declines steadily to -1.7 in the fifth; the low- h_c line is flat through the first two quintiles, dips sharply to -2.4 in the third, and partially recovers to -0.9 in the fourth. The two lines cross and show no systematic divergence, and the confidence intervals overlap at every quintile. The pattern is likely driven by the small and uneven cell sizes, which range from two observations in the lowest high- h_c quintile to nineteen in the highest. No low- h_c country falls into the top AI-readiness quintile, reflecting the strong positive correlation between cognitive skills and AI readiness documented in Section 5.1. The cell means are therefore highly sensitive to individual country values.

A.5 M3b bivariate scatter -supplement to §5.5

Figure A.7 plots the bivariate relationship for M3b, with separate linear fits by h_c group.




Figure A.7

Both regression lines slope gently upward, indicating that AI-ready countries tend to capture a slightly larger bottom-quintile income share, and the high- h_c line is steeper than the low- h_c line (qualitatively in the direction H3 predicts). The visual pattern does not, however, translate into a statistically significant interaction in the multivariate specification. The two lines are nearly parallel and the confidence intervals overlap throughout, and both groups are densely intermixed in the central range of AI readiness. The contrast with Figure 5.9 underscores the difference between M3 and M3b: in absolute poverty, the low- h_c gradient is visibly steeper and matches the regression, but in relative inequality, the gradient is shallow, ambiguous, and statistically indistinguishable from zero.

A.6 Robustness tables R4, R5, R16 - supplement to §5.6

Column (2) of Table A.2 re-estimates M2 after removing the 14 countries whose Cook's distance exceeds the conventional $4/N$ threshold.

R4: M2 + FE Full vs. Influential Excluded

Dependent variable: tfp_growth_avg

	M2 + FE Full (1)	M2 + FE Excl. Influential (2)
rc_c	-2.706 (3.244)	-4.216 (2.708)
hc_c	2.793 (2.743)	0.428 (2.450)
rc_hc	-2.327 (5.786)	-1.058 (4.886)

log_gdp_pc_2019	-0.836 (0.673)	-0.121 (0.526)
wgi_avg	-0.338 (0.772)	-0.519 (0.613)
trade_pct_gdp	-0.002 (0.006)	-0.005* (0.003)
inv_pct_gdp	-0.013 (0.060)	0.017 (0.046)
edu_spend_pct	0.085 (0.084)	0.033 (0.058)
pop_growth_pct	-0.143 (0.433)	-0.296 (0.336)
inflation_cpi	-0.027 (0.022)	-0.014 (0.095)
tfp_level_start	0.172 (13.330)	11.121 (10.928)
Constant	7.660 (15.657)	-9.154 (12.515)
World Bank region FE?	Yes	Yes
Observations	98	84
R ²	0.267	0.368
Adjusted R ²	0.111	0.228

Notes: *p<0.1; **p<0.05; ***p<0.01. Robust (HC3) SEs in parentheses below coefficients. Column (2) excludes 14 influential observations flagged by Cook's distance ($D > 4/N$).

Table A.2

The AI-readiness coefficient moves from -2.706 to -4.216, so the magnitude rises by roughly half as the most influential cases are removed, yet its standard error also tightens (from 3.244 to 2.708), so the coefficient remains statistically insignificant. The interaction term attenuates from -2.327 to -1.058 and stays insignificant. The cognitive-skills main effect collapses from 2.793 to 0.428, indicating that the bulk of the unconditional h_c -TFP-growth association was concentrated among the excluded economies. Model fit improves markedly (R^2 rises from 0.267 to 0.368, adjusted R^2 from 0.111 to 0.228). Trade openness becomes marginally significant at -0.005 (SE = 0.003, $p < 0.10$). The qualitative conclusion is unchanged: the interaction is indistinguishable from zero.

Column (2) of Table A.3 removes the 19 transition economies (former Soviet and Eastern European countries whose TFP series are sensitive to the post-1989 output collapse and the switch from Soviet-era to UN national-accounts standards (*Campos & Coricelli, 2002; Inklaar & Timmer, 2013*)).

R5: Excluding Transition Economies

Dependent variable: *tfp_growth_avg*

	M2 + FE Full (1)	M2 + FE No Transition (2)
rc_c	-2.706 (3.244)	-2.327 (3.547)
hc_c	2.793 (2.743)	4.942* (2.749)

rc_hc	-2.327 (5.786)	-1.375 (6.944)
log_gdp_pc_2019	-0.836 (0.673)	-0.990 (0.682)
wgi_avg	-0.338 (0.772)	-0.179 (0.941)
trade_pct_gdp	-0.002 (0.006)	-0.002 (0.007)
inv_pct_gdp	-0.013 (0.060)	0.001 (0.065)
edu_spend_pct	0.085 (0.084)	0.075 (0.091)
pop_growth_pct	-0.143 (0.433)	0.108 (0.645)
inflation_cpi	-0.027 (0.022)	-0.025 (0.021)
tfp_level_start	0.172 (13.330)	12.769 (13.537)
Constant	7.660 (15.657)	-4.156 (17.296)
World Bank region FE?	Yes	Yes
Observations	98	79
R ²	0.267	0.227
Adjusted R ²	0.111	0.012

Notes: *p<0.1; **p<0.05; ***p<0.01. Robust (HC3) SEs in parentheses below coefficients. Column (2) re-estimates M2 + FE on the sub-sample excluding post-socialist transition economies.

Table A.3

Sample size falls from 98 to 79. The AI-readiness coefficient is essentially unchanged at -2.327 (SE = 3.547), and the interaction attenuates to -1.375 (SE = 6.944); both remain far from significance. The cognitive-skills main effect, however, strengthens from 2.793 to 4.942 and crosses the 10% threshold (SE = 2.749, $p < 0.10$). The gain is consistent with the observation that many transition economies combine high Soviet-era cognitive scores with TFP trajectories shaped by structural transformation rather than by the skill-technology channel, so their presence attenuates the unconditional h_c -TFP-growth association in the full sample. Among the controls, population growth drops to near zero (0.108, SE = 0.645), while the initial-TFP-level coefficient remains large in magnitude but insignificant (12.769, SE = 13.537). Model fit deteriorates ($R^2 = 0.227$, adjusted $R^2 = 0.012$), as expected when the sample is restricted to a more homogeneous set of countries with less between-unit variation in the conditioning variables. The interaction retains its sign and remains insignificant.

Column (2) of Table A.4 re-estimates M2 on TFP growth computed over 2015-2019, retaining the 2022 GAIRI score as the measure of AI readiness.

R16: Placebo Test: M2 + FE on 2015-2019 TFP

Dependent variables: *tfp_growth_avg* (Main: 2022-2023) | *tfp_growth_placebo* (Placebo: 2015-2019)

	Main (2022-2023)	Placebo (2015-2019)
rc_c	-2.706 (3.244)	0.782 (1.899)

hc_c	2.793 (2.743)	2.160 (1.847)
rc_hc	-2.327 (5.786)	2.596 (5.038)
log_gdp_pc_2019	-0.836 (0.673)	-0.957*** (0.360)
wgi_avg	-0.338 (0.772)	0.166 (0.500)
trade_pct_gdp	-0.002 (0.006)	-0.002 (0.003)
inv_pct_gdp	-0.013 (0.060)	-0.057* (0.034)
edu_spend_pct	0.085 (0.084)	0.052 (0.047)
pop_growth_pct	-0.143 (0.433)	-0.017 (0.276)
inflation_cpi	-0.027 (0.022)	-0.019 (0.015)
tfp_level_start	0.172 (13.330)	2.049 (3.088)
Constant	7.660 (15.657)	8.168** (4.148)
World Bank region	Yes	Yes
FE?		
Observations	98	98
R ²	0.267	0.415
Adjusted R ²	0.111	0.291

Notes: *p<0.1; **p<0.05; ***p<0.01. Robust (HC3) SEs in parentheses below coefficients. Both columns use the same M2 + FE specification but differ in the dependent variable: column (1) uses TFP growth 2022–2023 (main outcome); column (2) uses TFP growth 2015–2019 (placebo period). rc is still measured in 2022 in both columns (lagged regressor), so a zero/weak association in the placebo column supports the main specification.

Table A.4

Sample size is unchanged at 98. The interaction coefficient is 2.596 (SE = 5.038), so opposite in sign to the baseline estimate of -2.327 and wholly insignificant. The AI-readiness main effect is positive but insignificant (0.782, SE = 1.899), in contrast to the negative but insignificant baseline value of -2.706. The placebo specification therefore produces no evidence that the AI-readiness × cognitive-skills interaction predicted TFP growth before AI technologies began diffusing at scale.

The placebo window does, however, reveal a tighter overall model. R² rises from 0.267 to 0.415 (adjusted R² from 0.111 to 0.291). Log initial GDP per capita is negative and highly significant (-0.957, SE = 0.360, p < 0.01), indicating robust conditional convergence in the pre-AI period, in contrast to its insignificant baseline value of -0.836. Investment enters with a marginally significant negative coefficient (-0.057, SE = 0.034, p < 0.10). The sharper estimates reflect the longer averaging horizon: TFP growth annualised over four years smooths the idiosyncratic year-to-year noise that dominates a one-year growth rate, reducing residual variance and tightening the standard errors. The fact that the standard macro controls perform better in the placebo window is therefore a measurement artefact of the longer horizon rather than evidence that the underlying relationships have weakened since 2022.

A.7 Within-country panel fixed-effects - supplement to §5.7

Table A.5 re-estimates the AI-readiness/TFP association using only within-country variation through country fixed effects, adding year fixed effects and time-varying governance in two specifications (FE1-FE3). The panel covers 88 country-year observations over 2022-2024, with standard errors clustered at the country level.

Panel Country-Fixed-Effects			
<i>Dependent variable: tfp_growth (annual panel)</i>			
	FE1	FE2	FE3
	Country FE	Two-Way FE	TW + Gov.
	(1)	(2)	(3)
rc_annual	-1.724** (0.711)	-1.364 (0.877)	-1.190 (0.822)
trade_pct_gdp	-0.131* (0.076)	-0.140* (0.078)	-0.156** (0.069)
inv_pct_gdp	-0.191 (0.191)	-0.239 (0.206)	-0.263 (0.166)
inflation_cpi	-0.183*** (0.018)	-0.181*** (0.019)	-0.170*** (0.013)
pop_growth_pct	-1.210* (0.714)	-1.045 (0.736)	-0.909 (0.730)
edu_spend_pct	0.152 (0.169)	0.160 (0.170)	0.125 (0.153)
wgi_avg			25.896*** (9.095)
Observations	88	88	88
R ²	0.775	0.776	0.823
Adjusted R ²	0.484	0.474	0.573

Notes: *p<0.1; **p<0.05; ***p<0.01. Cluster-robust SEs (clustered by country) in parentheses below coefficients. Column (1) includes country fixed effects only; column (2) adds year fixed effects (two-way FE); column (3) adds the World Bank governance average (wgi_avg) to the two-way specification. rc_annual is the annual within-country [0,1]-normalised GAIRI score. The hc channel (hc_c, rc × hc) is intentionally omitted in this panel specification.

Table A.5

The coefficient on AI readiness is negative and significant in FE1 (-1.724, SE = 0.711, p < 0.05), but loses significance once year fixed effects are added: -1.364 (SE = 0.877) in FE2 and -1.190 (SE = 0.822) in FE3. The widening of the standard errors between FE1 and FE2 reflects the loss of identifying variation once the common post-pandemic recovery year is partialled out. Adding time-varying governance in FE3 attenuates the coefficient further, indicating that part of the FE1 effect is entangled with year-to-year movements in institutional quality rather than with AI readiness per se. The attenuation across the three columns contrasts with the cross-sectional M2 estimate of -2.706 (SE = 3.244), which is statistically indistinguishable from zero and far less precise than the panel estimates.

Figure A.8 visualises the three within-country point estimates cluster between -1.7 and -1.2, with FE1 the only one whose 95 % confidence interval excludes zero, while the M2 coefficient sits near zero with a confidence interval many times wider. Coefficient magnitudes are not comparable across the two panels because of the different τ_c normalisations.




Figure A.8

Among the controls, inflation is the most consistent predictor across the panel, significant at the one per cent level throughout: -0.183 (SE = 0.018) in FE1, -0.181 (SE = 0.019) in FE2, and -0.170 (SE = 0.013) in FE3. Trade openness is negative and marginally significant in FE1 and FE2 (-0.131 and -0.140, both $p < 0.10$) and reaches the 5% level in FE3 (-0.156, SE = 0.069). Population growth enters with a marginally significant negative coefficient in FE1 (-1.210, SE = 0.714) but loses significance once year fixed effects are introduced. Investment is negative and education spending is positive in all three specifications, but neither approaches conventional thresholds. Adding time-varying governance in FE3 has two visible effects: the WGI composite itself enters with a large and highly significant positive coefficient (25.896, SE = 9.095, $p < 0.01$), and it absorbs part of the variation previously attributed to AI readiness, moving the AI-readiness coefficient from -1.724 in FE1 to -1.190 in FE3. The R^2 rises from 0.775 in FE1 to 0.776 in FE2 and 0.823 in FE3 (adjusted R^2 : 0.484, 0.474, 0.573 respectively).

Figure A.9 plots the demeaned data underlying FE1. Each point is a country-year observation with both the AI-readiness score and the TFP growth rate expressed as deviations from the country's own mean. The bivariate within-country fitted line is essentially flat, with large dispersion on both axes. Ukraine appears as a prominent outlier on both sides of the origin, reflecting the unusually large swings in measured TFP growth experienced during the war period. The within-country standard deviation of AI readiness is very low, with at most two to three annual observations per country. Year-to-year movements in the GAIRI index are small relative to the cross-country spread, so much of what the fixed-effects specifications identify reflects transient noise in both the index and annual TFP growth rather than a structural within-country relationship. The grey slope shown is the bivariate within-country association.




Figure A.9

The loss of significance in FE2 and FE3 indicates that the apparent within-country AI-readiness/TFP association in FE1 is partly absorbed by common time effects, leaving only inflation and (weakly) trade as robust within-country correlates of TFP growth in this panel. Together with the flat demeaned scatter, this suggests the cross-sectional null in M2 is not driven solely by omitted country characteristics, while remaining fragile to the inclusion of year fixed effects and of time-varying institutional quality. The within-country estimates may also reflect mean-reversion in TFP growth and reverse causality from macroeconomic conditions to the GAIRI index over a window this short.

A.8 Attribution statement

This work was primarily human-created. AI was used to make stylistic edits, such as changes to structure, wording, and clarity. AI was prompted for its contributions, or AI assistance was enabled. AI-generated content was reviewed and approved. The following model(s) or application(s) were used: ChatGPT, Claude, GLM.