

Generation AI Starts Before Consciousness

From Passive Exposure to Native Capability in the Beta Generation

The Economy Research Editorial^{1,2}

¹The Economy Research, 71 Lower Baggot Street, Dublin 2, Co. Dublin, D02 P593, Ireland

²Swiss Institute of Artificial Intelligence, Chaltenbodenstrasse 26, 8834 Schindellegi, Schwyz, Switzerland

Abstract

This paper extends recent analyses of AI exposure in early childhood by introducing a generational-economic framework. While existing literature emphasizes developmental risks and ethical considerations, this study argues that early, passive exposure to AI systems is producing a structurally distinct cohort—referred to as the Beta Generation—whose cognitive, behavioral, and labor-market adaptation costs will be significantly lower than any prior generation.

Building on Brookings' categorization of AI-enabled childhood environments, this paper reframes early AI exposure not merely as a developmental concern, but as a **pre-labor conditioning mechanism** that reshapes future human capital formation. The findings suggest that AI-native cohorts will not “learn AI” as a skill, but instead **internalize AI-mediated interaction as a default mode of cognition and production**, fundamentally altering skill hierarchies, employment structures, and institutional design.

1 Introduction - AI Is Not Adopted - It Is Imprinted

One strong intuition in technology policy is that adoption is a decision: a household purchases a device, a school acquires a software package, a worker learns a platform. The calque is voluntarist, and sequential technology appears; users reflect, institutions transform. But the texture of contemporary childhood contradicts that narrative. The worlds of early childhood within high-connectivity settings are, by now, thoroughly encased in algorithmic mediation, often through devices and services which, to a child and (more often than not) to a parent, do not identify themselves as "AI." A 2025 survey of US children 0–8 found that adult caregivers averaged 2 hours and 27 minutes of daily screen time, of which 2 hours and 58 minutes were consumed by children aged 5–8.^[1] This reliance pattern is not restricted to entertainment; the same survey documents growing mobile platform use in parents and caregivers as a co-actor of bedtime, dinner, and emotion regulation routines.^[2] In other words, "digital exposure" is not an afterthought inserted into children's daily lives but rather an extension of their daily lives.

An imperceptible but conclusive difference flows from time. The prior generations learned computers, then the internet, then platform ecosystems, then mobile. Even where adoption was prompt, that adoption was delayed relative to the emergence of the dominant language and social and cognitive tools. The generation that begins in the second half of the decade will emerge into home and school worlds in which generative systems and persistent suggestion engines are ubiquitous workhorses. The term "Generation Beta" refers to the demographic group following Generation Alpha, a name proposed by demographer Mark McCrindle for those born from 2025 onwards, though the exact range of years can vary depending on the definition used, according to McCrindle.^[3] The key point is not the defining cutoff but that the child's first environment can no longer be disentangled from artificial intelligence engines.

Within this framing, the common discussion in early childhood AI risk around content hazards and privacy is instructive but not comprehensive: focusing only on second-order developmental and institutional effects. What makes early insults unique is that they can shift baseline cognitive reasons for learning: what the learner perceives as knowing; what behavior gets reinforced; which behaviors receive the most attention, and how trust in human and machine systems is calibrated by discourse. The latter is significant because language and cognition develop in dynamic interactional microecologies; adult speech, child utterances, and turns are highly correlated in the speech of adult, child pairs. Longitudinal cohort studies of children aged 12 to 36 months, equipped with Language Environment Analysis technology, revealed that a 1-minute increase in screen time was predictive of lower scores on multiple measures of adult words, child vocalizations, and conversational turns at 36 months.^[4] This does not demonstrate that all digital usage is cognitively disruptive, or that there is no forthcoming causal path analysis. It merely shows that a screen-dominated interactional ecology can supplant the usual near-rikka interactional ecologies that nurture primary cognitive skills. Until similar effects are noted in AIian bedroom-based symbioses, the displacement risk is not primarily about its poison content, but about replacing contingency-rich human interaction with contingency-rich machine interaction.

A 2025 Pew Research Center report found that about 1 in 10 parents of children aged 5 to 12 say their child uses AI chatbots such as ChatGPT or Gemini.^[5] The report also notes that children in lower-income households

tend to spend significantly more time in front of screens compared to their peers in higher-income households.^[6]

The need for reframing is exorable, given the pace of AI diffusion from 2023 to 2026. In 2023, UNESCO warned of the risks of rapid adoption of generative AI in education, illustrating how quickly the technology integrates into core systems.^[7] Meanwhile, even as some countries reached 68% connectivity in 2024 and 74% in 2025 (per the International Telecommunication Union), gaps between high- and low-income sources persisted.^[8] According to a UNICEF report, about two-thirds of school-aged children worldwide, or 1.3 billion children between the ages of 3 and 17, do not have internet access at home. This lack of connectivity raises serious questions about labeling today's youth as a single, homogeneous "AI generation."^[9] Policymakers must both mitigate harm and recognize early AI exposure as a new input into human capital, with growing implications for inequality and institutional legitimacy.

2 From Tools to Environment: The Shift in Technological Ontology

The most consequential difference is ontological: AI is moving from a semi, optional toolkit to an ambient environment. In a toolkit paradigm, human agency is clear and within view: user, initiates, tool, responds, and the distinction between user and tool remains transparent. In an environment paradigm, AI has already begun designing the alternatives, guiding "attention, sorting facts," and providing system incentives, all without regard to whether users are cognizant of a technological intermediary. This is not a semantic difference; it is a governance distinction. Tool governance could be limited to single artifacts and interfaces, and provide user education. Ambient governance must confront infrastructures of recommendation, observation, and decision automation.

Three layers of reframing clarify how "AI as environment" shapes early childhood. The first layer, infrastructural AI, consists of systems that circulate around rather than with a child—monitoring, sorting, assigning amber flags, and discreetly optimizing interactions. The second layer, assisted cognition, includes systems that guide children through learning pathways, teach them to organize tasks, and shape their perceptions of difficulty, challenge, and competence. The third layer, social simulation, refers to relational assistants that converse, co-author stories, and serve as mental artifacts—first blending early childhood and computing, then later, design in general. These layers are not mutually exclusive; they may accumulate. A learning app is also a data collection device. A conversational agent is a profiling and advertising interface. A device in a smart, connected room can also act as a parent competing for the household's attention.

The structural explanation is that environments influence perception. If someone first learns the concept of "technology" within their culture, then they can recognize external devices as technology. If mediation is present before one's epistemological habits form, it is experienced as the normal way to access the world. This is central to the idea of "imprinting": how early one is introduced to mediation determines its internalization, rather than simply establishing a user-system relationship.

The macro drivers supporting this environment have intensified since 2023. Private investment in generative AI surged, making both consumer-facing and enterprise-facing deployment a strategic priority rather than an experimental extension. According to SP Global, private investment in generative AI grew to about 29 billion

dollars in 2023^[10] and surpassed 56 billion dollars in 2024, marking a significant increase and setting new records for the sector in both years.^[11] As investments grew, so did adoption: a 2025 AI trend policy synthesis found that the proportion of organizations using AI rose dramatically between 2023 and 2024.^[12] The implications for early childhood are subtle but substantial. As workplaces, schools, and public services adopted AI as an infrastructural layer, many families had to adjust how they planned their time and what they expected of each other, especially regarding AI-conscious productivity and coordination. Early childhood became a setting where families adapted to this new layer.

Environmental framing relates to order policy priorities as well. An AI as a "product" approach to policy ensures a child policy regime is predicated on age gating, child-safe interfaces, and consent. An AI as an "environment" approach to policy means setting systemic standards of accountability: procurement rules for public institutions, caps on profiling and behavioural ads, audit demand for recommendation engine, rights to non-participation that do not result in an educational penalty for children, In guidelines issues by UNICEF 2025 an apparent "product" focused regime,(stressing child friendly rules and controls that track developmental stages rather than treating childhood as some generic "user") contrasts with a "system" model, in which design of governance is "focused on designing systems that give children the possibility of participation without compromising the possibility of protection".^[13] That is an "environmental" thesis: children are not just users, they are developing persons, dynamic extenders of the socio-technical system.

3 Generational Skill Formation: From Adaptation Cost to Zero-Cost Integration

Generational analysis is often invoked imprecisely, as a condensed cultural category. Here, it serves a more rigorous purpose: pinpointing how the timing of an encounter modulates adaptation costs. 'Adaptation cost' is not only a matter of technical training epochs but also of cognitive frictions (architectural questions, system sets), epistemic frictions (knowing when and why to trust system outputs), and behavioral frictions (learning to recompose routines with new kinds of interfaces). Previous 'generations' bore hefty adaptation costs because they had to retrofit cognitive and institutional routines onto the interface from entirely unrelated frictions. The proposition here is that, for a child, formed minds exposed to an AI-saturated world, those retrofit costs diminish. The human capital in question is not "AI, competence" as a standalone skill, but interface fluency, delegation cues, and default prompt techniques learned by doing.

The empirical context strengthens the justification for a speculative moment: Many of the youngest learners will know AI and mediated learning functions, with tasks described by caregivers.^[5] Interestingly, generative AI 'started to go mainstream from months of its public release'a handful of months later.^[7] If the AI diffusion is stitched to the trajectory of childhood household device ownership, then the institutional prognostication is straightforward: By the time that 'entry cohort' of children are starting school, their educational experience will be structured around AI, mediated content/assessmentist incentive/cross-instructor hybrid pedagogies out of the box.

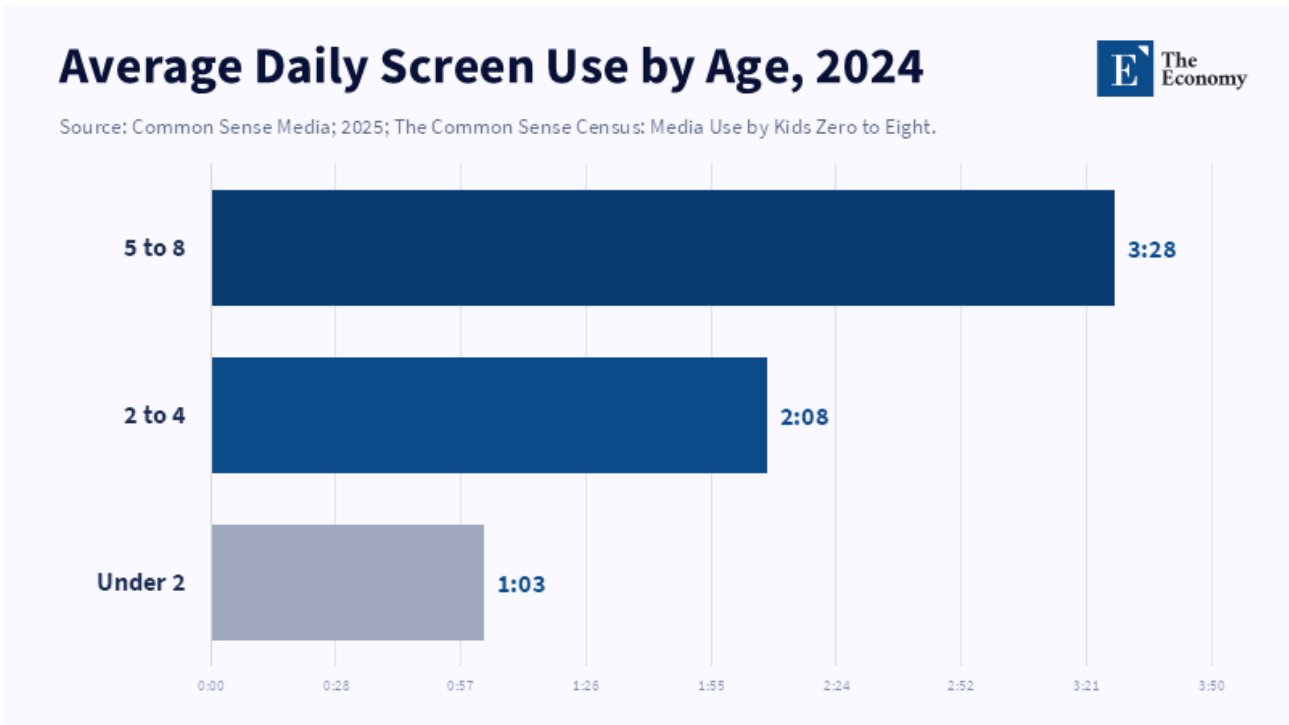


Figure 1: Early digital exposure already differs sharply even within childhood, suggesting that “AI-native” formation begins from unequal levels of mediated experience.

Examples from education systems exemplify how rapidly usage can ossify as normal. In 2024 (the first year that AIAP data was collected), the OECD’s Teaching and Learning International Survey (TALIS) asked teachers if they used AI for instruction or supporting students’ learning; among reporting OECD systems, the average level of teacher AI usage was high (although varying across countries), broken down as follows.^[14] In that same report, sizable percentages of teachers who used AI said it helped them learn about a topic efficiently, develop lesson plans, compose reports, or create assessments, but they also expressed reservations about academic integrity and bias amplification.^[15] The policy takeaway is that AI is entering the classroom not only through students but also through teachers’ workflows and expectations. With AI normalized in the preparation and assessment “environment,” children’s learning incentives shift: the marginal benefit of memorization or procedural reproduction drops, while the benefit of framing, selecting, and cross-verifying moves up.

Importantly, “zero-cost integration” is not an inclusive promise. Early exposure to digital media is prestratified. In the 2025 U.S. census of 0,8 media use, the gaps in screen time between households with very low and very high incomes were substantial.^[6] Since high levels of screen use are associated with less parent/child talk, the displacement effect is expected to be unequally distributed in ways that reinforce existing educational inequalities.^[4] Artificial intelligence-enabled platforms could, theoretically, offer compensatory scaffolding, including adaptive instruction, linguistically rich interaction, and specific feedback. But absent institutional protections, the more likely outcome is uneven substitution: resource-limited households will be less likely to invest in AI-based supervision of routines and attention management, while those with abundant resources will be best positioned to use AI for augmentation rather than substitution.




Figure 2: The growth of mediated environments does not translate uniformly into stronger reading habits, pointing to a gap between access and durable cognition.

4 The End of "Learning AI" as a Skill

The dominant policy approach to generative AI so far has focused on literacy: teaching students how to prompt, address model flaws, and use AI ethically. Those types of programs make perfect sense for generations learning to navigate and make use of AI-enhanced environments. They do not really make sense for generations where, as a matter of course, the world is already mediated by AI. For generations who were born and raised under AI, "learning AI" might quickly come to resemble another category mistake, in the same way as "learning search engines" would for a generation that leaves search as an ambient layer of knowledge. It seems more likely that what will matter at scale after a certain point will be touch interaction with generative systems becoming second nature, as contrasted with not being able to dominate the interface, the real game changer.

Labor market evidence already suggests as much. For instance, a 2024 analysis by the OECD on changes in demand for skills in AI-related occupations shows that four out of ten high-exposure AI job vacancies required management and business process skills across all OECD countries sampled in the study. Some of the types of skills that will be necessary for humans to work in tandem with AI have been identified. Another OECD 2024 report clearly confirms that a number of skills, for example, social, emotional, and digital skills, are in high demand in AI-intensive occupations, and that the demand for several types of skills increased over time.^[17] It finds, on its face, evidence supporting an optimistic outcome: AI will lead to a focus on human-centered skills. But it also finds evidence suggesting a different outcome: for instance, the implementation of AI might lead to declining demand for certain skills at the establishment level, such as cognitive and digital skills, resulting in less demand for these skills within the institution.^[18] The key qualitative insight, then, is to conceptualize this as a structural trade-off: while one can put policies and investments in place to drive productivity growth, AI could simultaneously reduce the institutional demand for certain skills.

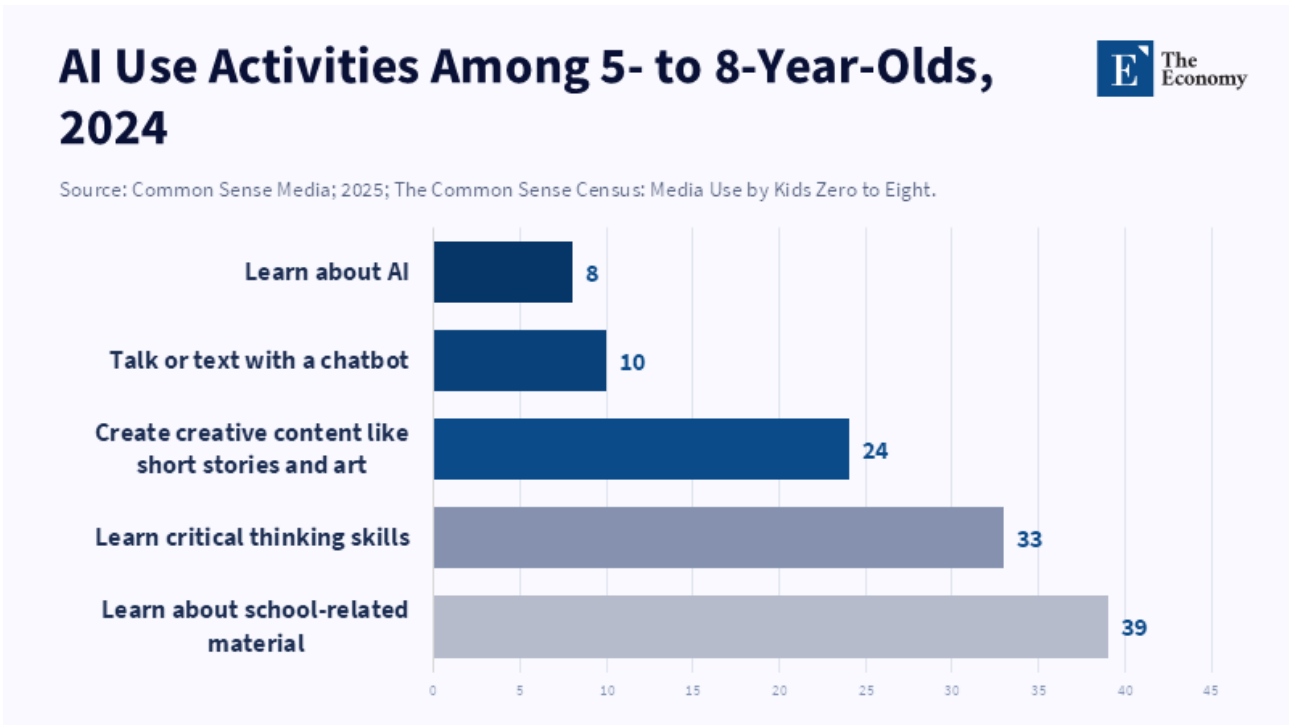


Figure 3: Native familiarity is not evenly distributed: the infrastructure of early digital conditioning is already shaped by household inequality.

Short-run productivity evidence also helps address the performance-development distinction. In a single lab experiment of professional writing tasks, a generative AI tool decreased time per task by 12,74% and increased output quality.^[20] In a real-world large field experiment in customer support, the availability of a generative AI conversational assistant increased productivity, with productivity increases exceeding 30%, especially for less experienced and lower-skilled workers.^[21] All of this, often cited in policy discussions, lends credibility to the productivity case for AI, offering support for the report’s boldest claim: that in the short run, by instantaneously enabling performance through production of scaffolding, AI will unchallengeably compete away superior ability, skewing institutions toward the development of scarcer differences.

This is the sense in which “AI literacy” must be recast for an AI, native cohort. No disclosure can be reduced to adherence to functioning rules. The ingrained capacities are Meta-cognitive and institutional: understanding when not to entrust, bolstering discretion and focus in the absence of machine kick, tellable, truthful text, and the safeguard of one’s own mental integrity in the face of centralizing culture. According to UNICEF, when introducing AI to children, it is important to encourage their resilience and independence, ensuring that AI is used to support children’s thinking rather than replace it.^[22]

Counterexamples need to be considered carefully, because they highlight where this story could be exaggerated. The most common example is the calculator: the argument goes that we have had technologies that automate cognitive tasks for a while, and education and labor markets have adapted to them. This comparison does not hold for three reasons. First, calculators are fixed in fixed referents and deterministic; generative systems are explorative and open-ended, and the outputs they generate can be especially persuasive in contextually sensitive ways that can obscure errors, increase, rather than reduce, the epistemic cost of validation. Second, calculators did not state that they would be conversational partners, while generative systems tend to. Third,

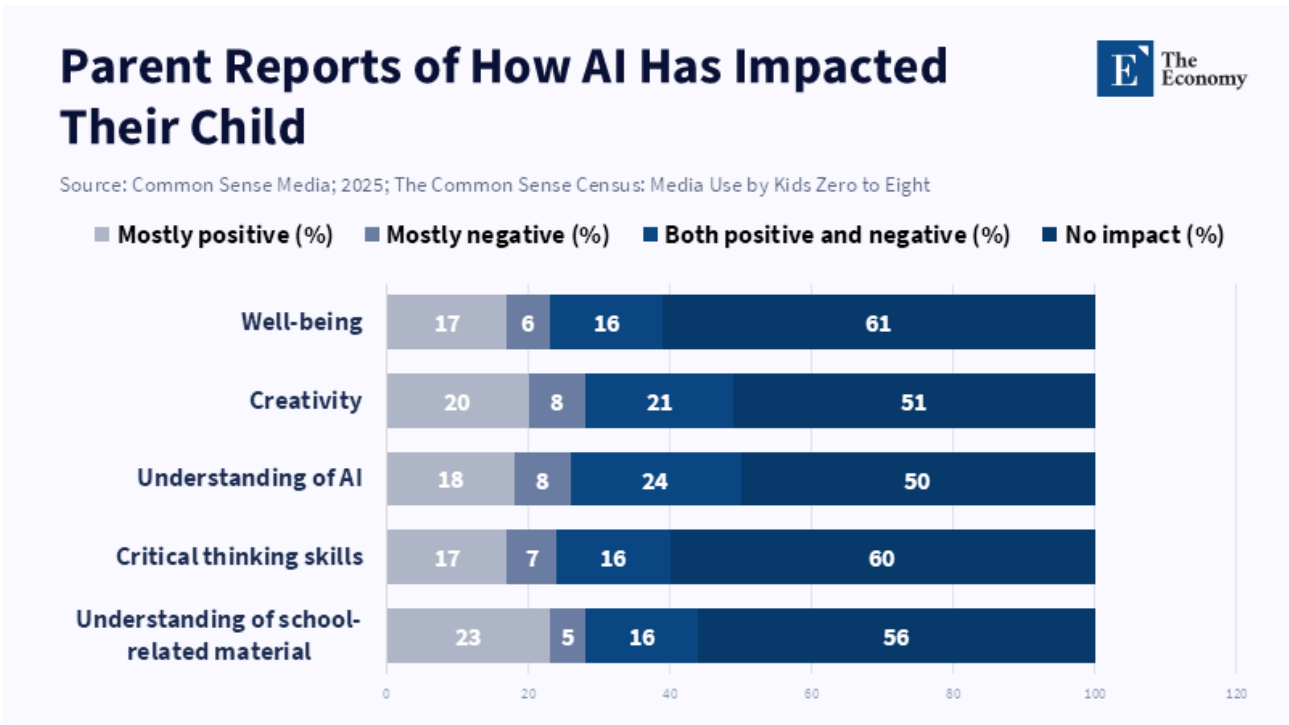


Figure 4: Children already use AI for practical tasks rather than for explicit “AI learning,” reinforcing the shift from skill acquisition to default usage.

generator modules are involved in data economies, increasingly fueling monetization, so the interface is hardly neutral even if the output is. UNESCO’s 2023 guidance recommends governance measures, including time-limited age restrictions for independent conversations with generative systems; educational decision-making must rely on the validation of it for pedagogic use.^[23] This is more than a formula for arithmetic; it is a claim that generalized systems will revolutionize assessment, motivation, and responsibility.

5 Labor Market Implications: The Collapse of Transitional Friction

Technological change in labor markets occurs as transitional friction: skills mismatch, retraining costs, and institutional lag. Much of the excitement about AI is based on a disruption story where a significant number of workers are disrupted more quickly than they can adapt. That scenario is not unlikely, but it errs by failing to account for the temporal heterogeneity of the workforce. Introducing generational differences in the costs of adaptation makes bifurcation more likely: an AI disruption may be localized to workers born in the pre-AI era, while new entrants are conditioned by AI as business as usual.

There are two strands of evidence for this two-pronged position. One is that automation risk remains pervasive across advanced economies. According to the OECD, automation risk is, on average, quite high across a sample of OECD countries, with the highest-risk occupations accounting for about 27 percent of employment.^[24] Additionally, a large number of discussion points point to the fact that such highly skilled occupational categories have relatively less automation risk, despite being more exposed to AI progress, because many highly skilled tasks contain bottlenecks that seem difficult to automate.^[25]

The second is that task-level studies of generative AI exposure suggest that the dominant impact will very often be augmentative rather than fully automating, with meaningful effects on work quality and work




Figure 5: Firms are adopting AI rapidly, meaning that the generation entering work will face a labour market where AI interaction is already normalized.

organization. An ILO working paper estimates that only some occupational categories are highly exposed to the task, level, and separate between their automation potential and their augmentation potential, also noting that the higher impact may largely result from augmentation effects with significant variation by country income and gendered exposure profiles.^[26]

Misunderstandings about the importance of augmentation are also common. If the effect of AI is to augment work, then the central labor market shock of the technology is less immediate job loss and more a new definition of what competence is required. Augmentation introduces new standards of what is expected to perform at work. For older cohorts, the costs of adapting to this new "standard" of what is performing at work are the costs of adaptation: new interaction protocols, new verification standards, and new task dimensions while simply maintaining output levels. For AI-native cohorts, those same protocols are no longer "new" but simply what is expected: the effect of AI on the labor market thus becomes intergenerational, while overall employment impacts remain more uncertain.

This recasts the familiar framing question about whether AI will "replace jobs" into a more policy-relevant question: who should bear the burden of transition? The answer will likely be stratified by age cohort, training received, and institutional setting. Where there are fast organizational step changes in AI integration, there is little room for people to derive economic value from procedural forms and an increased preponderance for governance over interfaces. Young cohorts brought up in these and other AI-enabled workplaces will have an easier time treating AI delegation as a routine matter. More senior cohorts will have ultimately been highly functional in domain expertise, but disproportionately suffer the additional friction of interface fluency, especially as working time devoted to training is cut back and the time to hands-on AI productivity is shortened.

The case for superhuman labor is built on the evidence of the AI revolution, without any overreach about



Figure 6: Firms are adopting AI rapidly, meaning that the generation entering work will face a labour market where AI interaction is already normalized.

innate intelligence” superiority. If superhuman is simply managing a more complex output with an AI-aided coordination, then the pool of candidates could get larger among cohorts for whom orchestrating with AI (like many of our instruments) is just a matter of habit, not an anchor of unsurvivable cost. Not that this will necessarily lead to such labor becoming ever more common; any high-end performance still hinges on specialized performance abilities that are independent of technological persistence, social intelligence, scruples, and subject-matter expertise. It does, however, remove some of the uncertainty of candidate pools in the future, and alert institutions to more clearly prune resources to scale up at potential demand, and to diagnose some of the messes in the landscape accordingly.

Policy responses follow naturally. Preparedness for the labor market cannot be reduced to “upskilling” in the narrow sense. It must encompass intergenerational mobility supports, unconditional income insurance, midcareer retraining, and workplace design that maintains autonomy and job quality in augmentation. The ILO also emphasizes the importance of job quality and work intensity (“work is intense” in the ILO’s terms) over headline employment effects in the integration of AI into work.^[27] For the Beta generation, the critical labor policy has a timeline before: to prevent interface fluency from replacing core competency, systems of education must sustain cognitively demanding work and show the graduating generation how to govern the results of their machines. Otherwise, labor markets will produce a population of fluent coders and scribblers unable to reason, verify, and be confident in their authority.

6 Human Relationship Substitution vs. Augmentation

Human relationship substitution is frequently considered the fundamental risk of AI in childhood: the replacement of caregivers, peers, and teachers by machinery. But that framing is over, determined. A more analytically

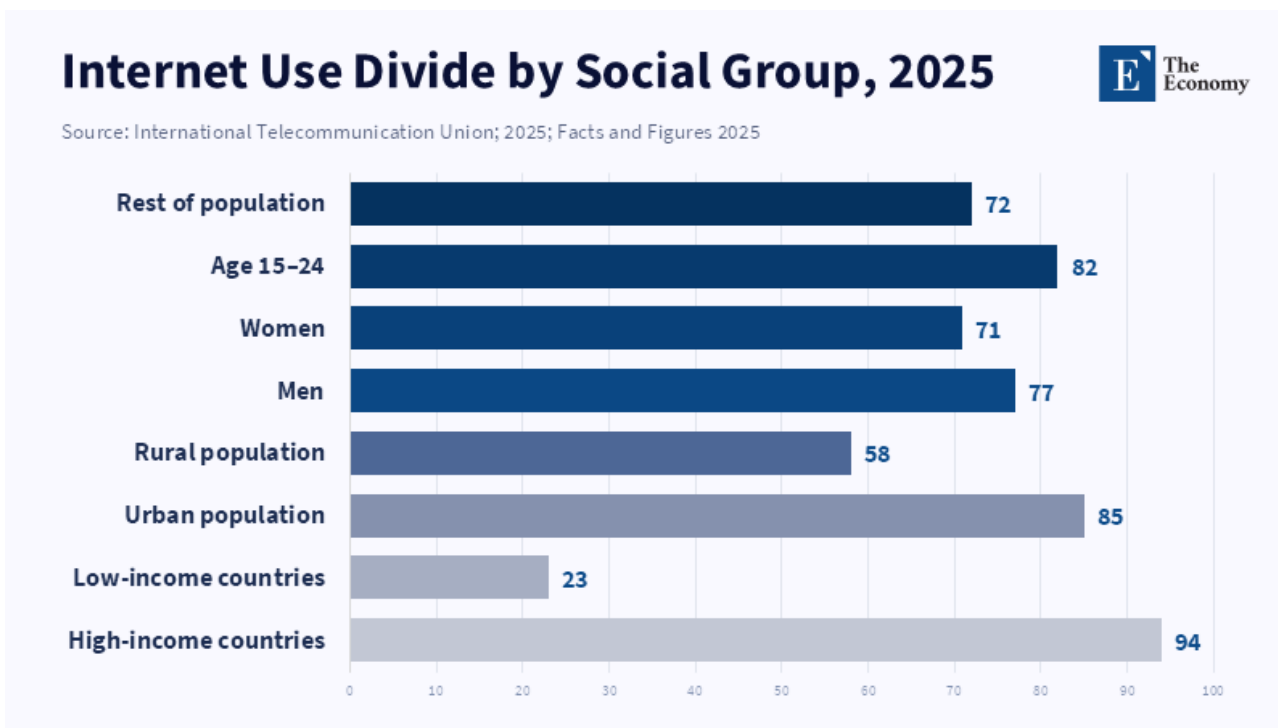


Figure 7: The coming AI-native workforce will not be socially uniform, because the connectivity base that supports digital fluency remains highly unequal.

productive distinction can be made between two equilibria: one that requires institutions and incentives to facilitate substitution, and the other that relies on them to facilitate augmentation.

There is evidence that children do not view conversational systems as neutral and disengaged tools. A mixed-methods, longitudinal investigation of children interacting with a smart speaker revealed that children overestimated systems’ mental and social capacities and had a limited understanding of the privacy and security implications.^[28] This is crucial because foundational social cognition is learned via attributions of agency, intention, and trustworthiness. If children regularly participate in systems that emulate the contingencies of conversative dialogue but have no embodied human limits, the developmental relevance is whether it is simply a replacement or an augmentation of the relationship?

Substitution, then, is not unavoidable and in principle not necessarily immoral. AI-mediated interaction could, in principle, support learning aims, accessibility, and creative interaction. Policy should not foreclose on those possibilities, in particular, for children with disabilities, or those who lack access to intensive expertise. But in the long run, the institutional danger is oversimplifying interaction into development; a system might deliver contingent responses that felt relational, but did not foster the social, cognitive skills that formerly arose from the mess of reciprocal human interactions.

This is where institutional bottlenecks come into focus. Educational systems may lag behind technological conditions, continuing to reward memory and procedural reproduction even in the face of generative AI that permanently collapses the cost structure for producing a baseline, acceptable surface-level output. According to the OECD’s 2023 review of generative AI governance, as of early 2024, the 18 countries and jurisdictions it examined had not designed any specific regulations for generative AI in education; governance mostly took the form of non-binding recommendations, and school-level decisions were disproportionately influential by.^[29] In

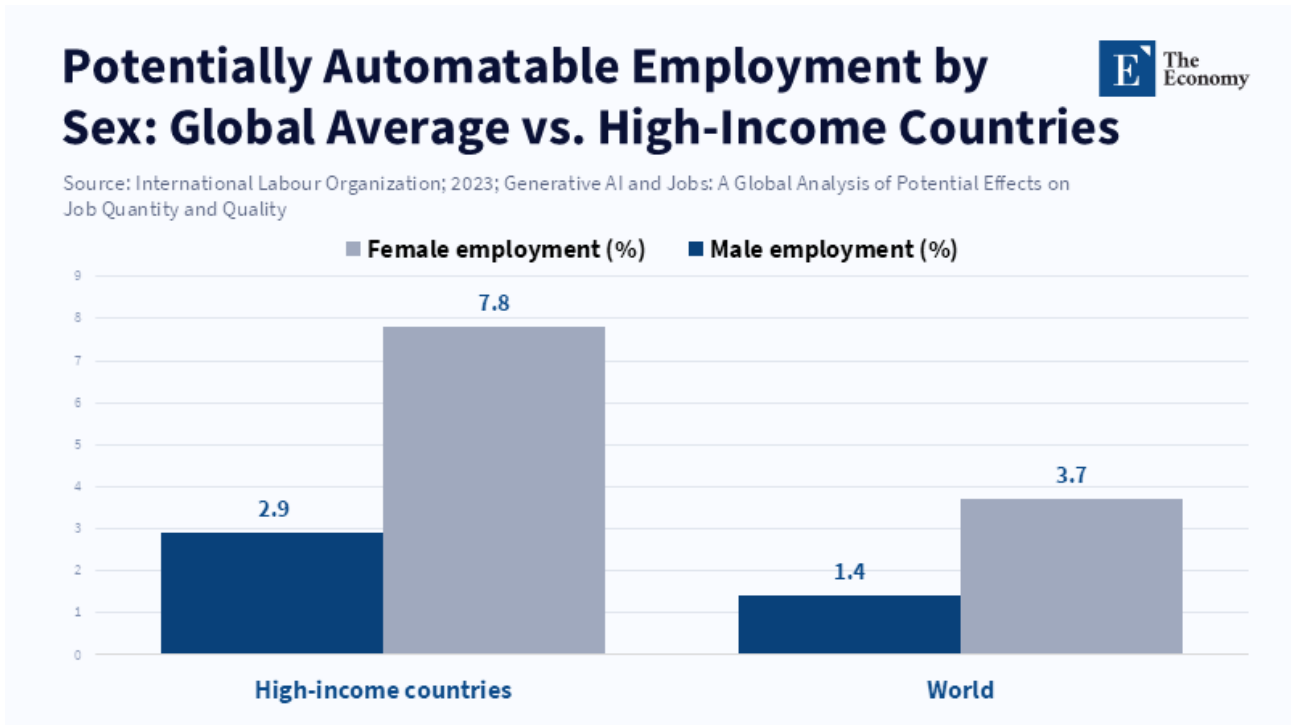


Figure 8: Reported effects are mixed rather than uniformly harmful or beneficial, which strengthens the case that institutions—not technology alone—will determine outcomes.

such a situation, substitution effects occur not because systems have been designed to serve a developmental function, but because institutions lack a coherent set of rules. Homework becomes unverifiable, testing becomes invested in game theory, and AI becomes a way to circumvent the institutional constraints rather than to transcend them.

The result is that the locus of danger moves from "AI itself" to institutional mismatch. According to a report by Steven Vosloo for the World Economic Forum, children are already using generative AI daily, including to help with homework, which may affect the amount of learning and critical thinking they engage in during important stages of development.[30] This is not meant as an argument against technology. It is a claim about institutional incentives. Where evaluation reinforces the formation of an answer, rather than understanding, AI offers an exhaust. Where instruction is reconceived as explanation, dialogue, and mastery of stabler mental operations, AI offers an aid.

Concrete implications follow on three levels of governance. For classroom users, curriculum design should foreground cognitive durability: prolonged reading, mathematical reasoning, and writing as thinking, not writing as styling. And AI should be a host for critique and revision, not Word's new black-box, first-pass reasoning. The goal is to maintain the developmental value of challenge, productive difficulty, and teach students to interrogate and validate machine thinking. For institutional users, institutional governance should move toward an integrated framing, in which ad hoc integrity warnings fall away: rules on AI use, repurposed assessment redesign, pedagogical use, guidelines for teachers, and procurement channels that enshrine privacy-by-design and auditability. And for modes of use of state/market policy: Child rights, governance of education, and labor, market strategy must meet. There are to be enforceable guarantees against profiling and exploitative data use involving children and facing systems, as well as investments that build teacher agency so AI isn't necessarily

a labor substitute for teachers in under-resourced schools.

The strongest pro-AI rebuttal is that AI democratizes access to expertise: A kid who received no tutoring can still have explanations tailored to him; a teacher who is drowning in paperwork still has more time for in-person time with some students. Those possibilities are quite real. But democratization is not guaranteed; it depends on whether we use AI as Infrastructure for the common good or merely as Attention Economy products optimized for engagement. Evidence that Households' screen exposure is already stratified by income cohort in the US by 2025 suggests that AI's benefits will likewise be stratified without policy solutions, if that infrastructure isn't specifically designed for equality.^[6] Those who seek the opportunity for augmentation must design for the "augmentation equilibrium," not hope for it.

7 Conclusion - From Child Development to Economic Infrastructure

Early experience with AI is no longer a marginal developmental variable. It is becoming a macroeconomic strategic variable whose effects will influence future labor supply, shape the burden of adaptation costs, and determine the institutional relevance of schools and training systems. The primary argument is not that AI disadvantages children per se, nor that a singular "AI generation" worldwide will have a homogeneously AI-exposed childhood. It is that AI will be used to mediate with children earlier, not later, producing a cohort for whom being mediated by AI is normal, and greatly reducing the friction in adapting to AI-saturated workplaces and education systems. Policy in this world is no longer concerned with the question of whether AI makes children worse off, but instead should be: what kind of economic agents does AI exposure before consciousness shape, and to what extent? Early evidence points to high early screen exposure, stark socio-economic gradients in media time, quantifiable reductions in parent and child speech, and high reported AI use in early childhood [1–4] and learning ^[5, 14]. Simultaneously, AI diffusion through institutions has accelerated the growth of large-scale organizational use, and increasing investment is further turning AI into infrastructure rather than a mere instrument ^[7, 12]. The key danger is institutional lag. By designing rules of work and study for a pre-AI era, we will assume flawed indicators of competence, misallocate investment in training, and drive deepening inequality. An adequate response demands governance commensurate with the upcoming structural transition, child-centric AI management, learning, capital, preservation of assessment changes, and high-adoption, cost-transition policies. Only thus can we ensure institutionally durable structures for an AI, native human capital cohort able to work and learn in a profoundly different economic landscape.

References

- [1] Common Sense Media (2025) *The 2025 Common Sense Census: Media Use by Kids Zero to Eight*. San Francisco, CA: Common Sense Media.
- [2] Common Sense Media (2025) *The 2025 Common Sense Census: Media Use by Kids Zero to Eight*. San Francisco, CA: Common Sense Media.

- [3] Caballero, J. and Fengler, M. (2025) ‘4 things to know about Generation Beta and how they could affect the global economy’, *World Economic Forum*, 18 February.
- [4] Brushe, M.E., Haag, D.G., Melhuish, E.C., Reilly, S. and Gregory, T. (2024) ‘Screen Time and Parent-Child Talk When Children Are Aged 12 to 36 Months’, *JAMA Pediatrics*, 178(4), pp. 369–375. doi:10.1001/jamapediatrics.2023.6790.
- [5] Common Sense Media (2025) *The 2025 Common Sense Census: Media Use by Kids Zero to Eight*. San Francisco, CA: Common Sense Media.
- [6] Miao, F. and Holmes, W. (2023) *Guidance for Generative AI in Education and Research*. Paris: UNESCO.
- [7] International Telecommunication Union (2024) *Measuring Digital Development: Facts and Figures 2024*. Geneva: International Telecommunication Union.
- [8] UNICEF Innocenti – Global Office of Research and Foresight (2025) *Childhood in a Digital World: Screen Time, Digital Skills and Mental Health*. Florence: UNICEF Innocenti.
- [9] Stanford Institute for Human-Centered Artificial Intelligence (2024) *Artificial Intelligence Index Report 2024: Chapter 4 – Economy*. Stanford, CA: Stanford University.
- [10] Stanford Institute for Human-Centered Artificial Intelligence (2025) ‘Economy’, in *The 2025 AI Index Report*. Stanford, CA: Stanford University.
- [11] Stanford Institute for Human-Centered Artificial Intelligence (2025) *AI Index Report 2025: Policy Highlights*. Stanford, CA: Stanford University.
- [12] UNICEF Innocenti – Global Office of Research and Foresight (2025) *Guidance on AI and Children: Version 3.0 – Recommendations for AI Policies and Systems That Uphold Child Rights*. Florence: UNICEF Innocenti.
- [13] Borgonovi, F., Bastagli, F., Ochojska, M. and Piumatti, G. (2025) ‘AI adoption in the education system: International insights and policy considerations for Italy’, *OECD Artificial Intelligence Papers*, No. 52. Paris and Turin: OECD Publishing and Fondazione Agnelli. doi:10.1787/69bd0a4a-en.
- [14] International Telecommunication Union (2025) *Measuring Digital Development: Facts and Figures 2025*. Geneva: International Telecommunication Union.
- [15] Green, A. (2024) ‘Artificial intelligence and the changing demand for skills in the labour market’, *OECD Artificial Intelligence Papers*, No. 14. Paris: OECD Publishing. doi:10.1787/88684e36-en.
- [16] Green, A. (2024) ‘Artificial intelligence and the changing demand for skills in the labour market’, *OECD Artificial Intelligence Papers*, No. 14. Paris: OECD Publishing. doi:10.1787/88684e36-en.
- [17] Noy, S. and Zhang, W. (2023) ‘Experimental evidence on the productivity effects of generative artificial intelligence’, *Science*, 381(6654), pp. 187–192. doi:10.1126/science.adh2586.

- [18] Brynjolfsson, E., Li, D. and Raymond, L. (2025) ‘Generative AI at Work’, *The Quarterly Journal of Economics*, 140(2), pp. 889–942. doi:10.1093/qje/qjae044.
- [19] UNICEF Innocenti – Global Office of Research and Foresight (2025) *Guidance on AI and Children: Version 3.0 – Recommendations for AI Policies and Systems That Uphold Child Rights*. Florence: UNICEF Innocenti.
- [20] OECD (2023) *OECD Employment Outlook 2023: Artificial Intelligence and the Labour Market*. Paris: OECD Publishing. doi:10.1787/08785bba-en.
- [21] Gmyrek, P., Berg, J. and Bescond, D. (2023) *Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality*. ILO Working Paper No. 96. Geneva: International Labour Organization. doi:10.54394/FHEM8239.
- [22] Andries, V. and Robertson, J. (2023) ‘Alexa doesn’t have that many feelings: Children’s understanding of AI through interactions with smart speakers in their homes’, *Computers and Education: Artificial Intelligence*, 5, article 100176. doi:10.1016/j.caeai.2023.100176.
- [23] Vidal, Q., Vincent-Lancrin, S. and Yun, H. (2023) ‘Emerging governance of generative AI in education’, in *OECD Digital Education Outlook 2023: Towards an Effective Digital Education Ecosystem*. Paris: OECD Publishing. doi:10.1787/c74f03de-en.
- [24] UNICEF Innocenti – Global Office of Research and Foresight (2025) *Guidance on AI and Children: Version 3.0 – Recommendations for AI Policies and Systems That Uphold Child Rights*. Florence: UNICEF Innocenti.
- [25] OECD (2023) *OECD Employment Outlook 2023: Artificial Intelligence and the Labour Market*. Paris: OECD Publishing. doi:10.1787/08785bba-en.
- [26] Gmyrek, P., Berg, J. and Bescond, D. (2023) *Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality*. ILO Working Paper No. 96. Geneva: International Labour Organization. doi:10.54394/FHEM8239.
- [27] Gmyrek, P., Berg, J. and Bescond, D. (2023) *Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality*. ILO Working Paper No. 96. Geneva: International Labour Organization. doi:10.54394/FHEM8239.
- [28] Andries, V. and Robertson, J. (2023) ‘Alexa doesn’t have that many feelings: Children’s understanding of AI through interactions with smart speakers in their homes’, *Computers and Education: Artificial Intelligence*, 5, article 100176. doi:10.1016/j.caeai.2023.100176.
- [29] Vidal, Q., Vincent-Lancrin, S. and Yun, H. (2023) ‘Emerging governance of generative AI in education’, in *OECD Digital Education Outlook 2023: Towards an Effective Digital Education Ecosystem*. Paris: OECD Publishing. doi:10.1787/c74f03de-en.

- [30] UNICEF Innocenti – Global Office of Research and Foresight (2025) *Guidance on AI and Children: Version 3.0 – Recommendations for AI Policies and Systems That Uphold Child Rights*. Florence: UNICEF Innocenti.