


We Teach Driving. Every Child Now Needs to Ride.

Education in the age of AI, and why every child must learn to create.

FIROZ AZEES, ILAN000



COVER · THE HORSE IS AI, WRITTEN IN ITS OWN WORDS

CONTENTS

PROLOGUE

A room in Rome

01

Education passes on two halves

02

Both halves were learned by living

03

School copied the production line and took the living out

04

The economy shifted, and AI removed the reason

05

From driving to riding

06

Creating brings both halves back

07

Every child gets to ride

08

Careers

09

What we know, and what goes beyond it

ABOUT
ILAN000

What Ilanooo does today

APPENDIX A

Sources and notes

APPENDIX B

How the Direction Index is built

In one page

Education passes on two halves. Half one is the tools we made to understand the physical world: language, numerals, mathematics, science. What they describe is there to be discovered. Nothing here looks down on these tools; what has to change is how they are handed on. Half two is the structure we made so that strangers can work together: money, law, status, culture. We make it. Both are compression, and both have to reach every child.

For most of history children learned both by living, by taking part in the work around them. Then school copied production: scientific, standardised, built on economies of scale, with one curriculum, one pace and one test. What fits through a production line is what one person can say aloud and many can repeat back, so school kept knowing and understanding and took the living out. Half two was never given a place on the timetable, and where it has come back as a subject, it has come back small.

The economy school copied is changing, and the machine now does much of the procedure. AI removes the reason for the line. Interaction that adapts to each learner, which the standardising was built to avoid, can now be had at a price. Employers expect two in five core skills to change by 2030.

AI is a horse. School teaches driving: set practices, passed on and checked. That method does not fit a horse, which has habits, changes in steps, and has to be learned by riding. A child who hands the job over and takes what comes back is generating. A child who directs the machine is creating. The difference decides what a child keeps. In one Turkish school, students given a version that simply answered scored 17% lower on the exam than students who never had it.

Creating sits at the top of the school's own list of what learning is for, and it brings both halves back in the same afternoon. Both halves should be learned by creating, and AI can make that possible for every child, not for a few. It builds the feel for the horse. Nobody has tested this yet with children and AI.

This is about equal access, not about closing schools. If the strongest models reach only some children, the gap widens. I want every child to get a lot of tokens through an adult institution, and I want us to watch what they do with them. Run it as a trial, with one group given tokens only, one given tokens and guided creating, and one given neither. Measure what is made, how much of it the child directed, and, once somebody builds the instrument, what the child understands about money, audience and trust.

We ask governments and funders to pay for the trial and the adults, AI providers to publish aggregate use by country and age band, and schools, libraries and community centres to be the place.

A room in Rome

On 6 January 1907, at 58 Via dei Marsi, in the San Lorenzo quarter of Rome, a ground-floor room in a tenement opened its door to about fifty children.

They were between three and six years old. San Lorenzo was a poor quarter and the parents went out to work, so the children raised themselves on the stairs and in the courtyard. The owners of the building, a group of bankers with more than four hundred tenements in the city, had a practical complaint. The children kept scratching the new whitewash off the walls. The engineer who ran the company, Edoardo Talamo, asked a physician named Maria Montessori to do something about it.

She had almost no money. Her biographer Rita Kramer writes that the owners were "not prepared to spend a penny on toys or equipment." She had an empty room, some materials she had designed, and one adult: Candida Nuccitelli, about forty, the building porter's daughter, with no training at all. Montessori called it the Casa dei Bambini. The Children's House. Not a school.

What happened next is in her own book and in Kramer's. The children worked with the materials. Writing came first, about ten months after the room opened. Reading followed within days. Children who had been defacing walls were writing by that autumn.

Talamo's plan was a school in every house the company owned. Five years later Montessori published her list of the houses that followed. By the end of 1908 there were four, two of them in San Lorenzo. Then the sponsor and the doctor fell out over the publicity, and, in her own words, the porter was told not to let her into the building.

The room is still there. A state school runs at 58 Via dei Marsi today. The method that started in it now gives its name to nearly sixteen thousand schools in 154 countries, by one 2022 census. About nine in a hundred of them are fully funded by a government.

Reformers and parents have been trying to repeat that room for more than a hundred years. For most children it has stayed an exception. I want to know why, and what it would take to stop it being one.

58 VIA DEI MARSI · SAN LORENZO · ROME

ABOUT FIFTY CHILDREN, AGED THREE TO SIX

6 JANUARY 1907




FIG. 1 · THE ROOM AT 58 VIA DEI MARSI, 6 JANUARY 1907

Education passes on two halves


Every generation hands the next one two things.

A child born today does not start where the first humans started. She starts where we are. Nobody could cover that distance in one life. It takes a way of handing on.

What gets handed on falls into two halves.

Half one is the tools we made to understand the physical world and move through it. Language. Numerals. Mathematics. Science. The physical world does not care what we think about it. It is there to be discovered, and what we find stays found.

Half two is the structure we made so that strangers can still work together. Money. Law. Status. The stories a culture tells about itself, and the names it gives to what we feel. We did not discover any of it. We made it to cooperate, because cooperation is how a weak animal survived, and cooperation at that size needs rules that millions of heads carry the same way.



BOTH ARE COMPRESSION. BOTH HAVE TO REACH EVERY CHILD.

FIG. 2 · WHAT EDUCATION PASSES ON

Both halves are compression.

Language lets one person pass on in an afternoon what another took a lifetime to learn. Numerals did it for counting. Before them you counted by repeating the thing: fish, fish, fish, fish. Then numerals came along, and zero with them, and a whole way of keeping accounts got shorter. In 1299 the money-changers' guild of Florence ordered its members to write the sums in their account books in words, not in the new figures, and the new way took generations to win. In Kerala, where I am from, Madhava and his school had infinite series for sine and cosine by about 1400, more than two centuries before Newton and Leibniz. Each of these cut the cost of passing something on. What took a generation to learn could be learned in a school year.

Half two works the same way. Money squeezes what a whole city owes and trusts into a number you can carry in a pocket. A contract is a handshake that has to outlive the two who shook hands. In 2025 Farrell, Gopnik, Shalizi and Evans, writing in *Science*, observed that markets and bureaucracies already work by lossy summary. A price, they write, "compresses vast swathes of detail into a simplified but usable representation."

Nothing here looks down on half one. A child who cannot read or count is shut out of everything else, and these tools are useful. What has to change is how they are handed on.

Both halves were learned by living

Where we can see it, a child's life and a child's learning were the same thing.

No child was ever taught grammar before she spoke. She was spoken to all day by the adults around her, who needed something from her, and she answered. Nobody planned a sequence. By four she was running a language that nobody had sat down and taught her. Language, the first of our tools, was learned by living.

So was most of the rest. In societies where children still grow up inside the family's work, anthropologists report the same pattern: children learn by watching and helping, with very little instruction. Barbara Rogoff calls it learning by observing and pitching in, and describes it in Indigenous-heritage communities of the Americas. David Lancy, reading the Six Cultures study of the 1960s, notes that only one of the six communities did not routinely put its children to chores. It was a middle-class town in the United States.

Those are societies of the present. Nobody measured how children learned in the year 1400. What the records of the past show points the same way.

English apprenticeship contracts say little about the craft. "The details of what was to be learned were not specified," writes the historian Patrick Wallis. On conduct they are exact. The apprentice promised to obey, to stay out of taverns and card games, to work hard, and to "keep his master's secrets." A boy bound to a master was learning a trade. He was also being taken into a household, with its discipline, its discretion and its place in the town.

Apprenticeship was the exception. In English towns in 1600 and 1700 about one teenage boy in ten started one. Most children never did. They learned at home, beside the work, and no contract recorded it.

That is how half two was learned. Nobody had to teach a child that a promise costs something, or whom to trust with a secret, or what a day's work was worth to the person who paid for it. She learned it the way she learned to speak, by being inside it. Half two was never on a syllabus, because it never needed to be.

The evidence that it still works this way is narrow, and all of it points one direction. Swedish adoption records let economists separate what a child inherits from what she is raised in. Lindquist, Sol and Van Praag found that being raised by entrepreneurs raised a child's chance of becoming one about twice as much as being born to them. In American data, Fairlie and Robb found that business owners who had worked in a relative's business did much better than those who had not, while merely having a relative who owned one made little difference.

That is one slice of half two, the slice somebody has measured. Exposure behaves the same way. Bell, Chetty, Jaravel, Petkova and Van Reenen followed more than a million American inventors through patent, tax and school records. Children who moved when young to places with more inventors became inventors at higher rates.

I go further than the data. Much of what we call intelligence is exposure.

That was the arrangement. A child lived inside the work of the adults around her and met both halves without anyone dividing them. Then something came along that could teach a great many children at once.

BEFORE: THE DAY AS A LINE OF REAL THINGS



AFTER: THE SAME DAY AS A TIMETABLE



EQUAL BLOCKS. A BELL BETWEEN EACH ONE.

FIG. 3 · THE SAME EIGHT HOURS, BEFORE AND AFTER

School copied the production line and took the living out

School borrowed the logic of production. What that logic could not carry, it left outside the door.

School helped carry reading from about one adult in ten in 1820 to almost nine in ten today, by the count at Our World in Data. It was not the only teacher, and in some places, Sweden among them, reading was taught at home and in church first. I care about the shape school took, and what that shape cost.

Historians still argue about why states built schools: to make citizens, to keep order, to make workers, to spread a religion. I am not going to settle that. The shape of the classroom has a source, and the source is written down.

The age-graded classroom came first. Boston's Quincy Grammar School, in 1848, is usually credited as the first American school built around one grade, one room and one teacher: twelve classrooms of about fifty-six pupils. By the 1910s the shape had a language. Between 1911 and 1925, Raymond Callahan writes, "efficiency" procedures were applied to classroom learning, to teachers, to the program of studies, to the organization of the schools, and to whole school systems. A newspaper editorial of 1912, quoted by one of the leading school administrators of the day, said schools should measure their results the way Frederick Taylor's scientific management had cut the motions in laying brick and the fatigue in handling pig iron.

In 1916 that administrator, Ellwood Cubberley, wrote the sentence that says it plainly:

"Our schools are, in a sense, factories in which the raw products (children) are to be shaped and fashioned into products to meet the various demands of life. The specifications for manufacturing come from the demands of twentieth-century civilization, and it is the business of the school to build its pupils according to the specifications laid down."

A critic of the same years had already named the trade. "We have consented to measure the results of educational efforts in terms of price and product," he wrote, "the terms that prevail in the factory and the department store."

Callahan's conclusion is the part worth keeping. Nobody sat down and designed it. Administrators were under attack from businessmen who wanted schools run the way they ran their own firms, and they gave in. Tyack, who studied the same ground, has the educators building it themselves, borrowing the corporation's way of organising. Either way, the model they reached for came from production, and school copied it without asking whether learning works that way.

The logic is economies of scale: standardise the product, and the cost of each unit falls. In school the product was a child and the standard was one curriculum for all. That is a matter of how the work is organised, not of how many children one teacher has.

A production line works on what can be specified in advance. The same part, the same pace, the same check at the end. In school the part was the lesson, the pace was the timetable and the check was the test. What fits through that line is what one person can say aloud and many can repeat back, the lowest rungs of what school itself says learning is: knowing and understanding.

What does not fit is everything that depends on the day. A child learning to deal with someone who disagrees with her. A promise that costs something when it breaks. Something she made that another child wants and she does not want to give. The odd, unplanned, consequential part of a day is exactly what a production line is built to remove.

School became standardised and choreographed, and sterilised as well. Every hour was planned before the child arrived, and nothing in the day could go wrong in a way that cost her anything. The living did not come in a standard size, so it went.

A PRODUCTION LINE



A CLASSROOM



WHAT FITS THROUGH

What one person can say aloud
and many can repeat back.

Knowing.
Understanding.

WHAT DOES NOT

A disagreement to work through.
A promise that costs something.
Something made that another child wants.
The odd, unplanned, consequential day.

FIG. 4 · THE LINE AND THE CLASSROOM

Half one paid first. Language, which every child had learned by living, now arrived as exercises to be explained from the front of a room. Half two paid more. It never had a place on the timetable, and when the living went out of the day nothing took its place. Where it has come back as a subject, it has come back small. For most of a child's life it was left to be picked up wherever she happened to be.

Where the living survived, in a family business or a house full of talk, a child still got it. Where it did not, she got whatever came along.

School was built for a production economy, and it worked for one. The teachers did what it asked.

The economy shifted, and AI removed the reason

The economy school copied is changing, and AI removes the reason school was built the way it was.


Start with what employers say. The 2025 Future of Jobs report from the forum at Davos surveyed more than a thousand employers covering over fourteen million workers. They expect 39% of workers' core skills to change by 2030. That is lower than the 44% they gave two years earlier, and it is still two in five. The most sought-after core skill is analytical thinking, which seven in ten companies call essential, followed by resilience, flexibility and agility. The roles they expect to shrink most in absolute numbers are clerical and secretarial: cashiers and ticket clerks first, then administrative assistants and executive secretaries.

The US Bureau of Labor Statistics sees the same pattern. In its latest projections, for 2025 to 2035, typists fall by about 34%, data entry keyers by about a quarter, and cashiers by around 200,000 jobs. The edition before it said that "the use of automated systems, including AI, is expected to contribute to declining employment of office and administrative support workers."

These are the jobs of following a procedure and keeping a record. The machine is better at that than it was, and employers are asking for something else.

I read this as larger than a list of jobs. For two centuries economists counted land, capital, labour and, later, entrepreneurship as the things that produce. Capital was the tool. Labour was the person who used it. A machine that can draft, calculate, schedule and write is capital that does some of the using itself. In the parts of work that are procedure, labour and capital have started to merge. What stays with the person is the part that is judgment: deciding where the machine goes, and for whom.


BEFORE



LABOUR USES CAPITAL

AFTER

decides where it goes, and for whom



THE MACHINE DOES THE PROCEDURE. THE PERSON DECIDES.

FIG. 5 · WHO DOES THE USING, BEFORE AND AFTER

Half two follows production. Status, credentials, hierarchies, what a job is and what a person is for were all made to fit the way we produced things. When production changes, we remake them. A child who spends fifteen years learning the old arrangements and then walks into the rebuilt ones pays for it.

The remaking runs toward smaller, decentralised units. Work that needed a company can sometimes be done by one person with a machine, and a skill can be put to use for others without anyone founding anything.

The other change matters most for school.

The production line had a reason. A production economy needed standard output at scale, and economies of scale reward standardising: one curriculum, one pace, one test. Learning that adapts to each child, that follows what she does next and changes the task when she is ready, does not fit a line. That is why school choreographed every hour, and why the organic, purposeful interaction children once had inside the family's work was cut out.

AI removes that reason. The adaptive part of an adult's attention, explanation, practice and feedback that fit one learner, can now be had at a price, so the standardising no longer has to bind, and the interaction can be unchoreographed again: purposeful talk, argument, making things. Two trials show that attention that adapts to one learner is feasible. At Harvard, 194 physics undergraduates each had one lesson with an AI tutor and one with in-class active learning. They gained more than twice as much from the tutor lesson, in a median of 49 minutes against 60 (Kestin and colleagues, 2025). In Edo State, Nigeria, over 600 secondary students spent six weeks after school with a chatbot tutor and gained about a third of a standard deviation more than students without one. That is a World Bank working paper, not yet peer reviewed.

Neither is a trial of children making things. The Harvard students were adults at a university. What the two show is that the adaptive part is feasible. On our estimate, a child using AI for about a hundred thousand tokens a day would cost between roughly fifty cents and thirteen dollars a month in tokens, across small and mid-sized models at current list prices. Adults, devices and connectivity are not in that number, and nobody has priced them.

The other part, an adult who knows the child and helps decide what is worth making, is still needed, and nobody has priced it. A tutor that explains the same lesson to one child at a time is only the line made faster. With the adaptive part cheap, the adult's time can go to what the line cut out, the unplanned and consequential day. The reason for the line is gone.

From driving to riding

Education was built to teach driving. AI is a horse.

Children are already on it. In the United States, 64% of teens say they use an AI chatbot, about three in ten every day, and 54% have used one for schoolwork (Pew Research Center, autumn 2025, 1,458 teens). In Britain, 56% of children aged 8 to 17 have used AI, and 66% of those aged 16 and 17 (Ofcom, 2026, 3,426 children). For India we found no national figure.

Schools have answered in two ways. Some ban and police. On 11 May 2026 the Princeton faculty voted, with one dissent, to put proctors in every in-person exam from 1 July, ending a ban on proctoring that had stood for 133 years. In the university's own senior survey of over 500 students, 29.9% said they had cheated and 0.4% said they had reported a peer. Others teach AI as a subject. UNESCO's 2024 framework for students sets out twelve competencies in four areas, and the United Arab Emirates has taught AI as a formal subject in public schools, from kindergarten to Grade 12, since 2025. Banning treats the horse as a threat; teaching AI as a subject treats it as a car.

A driving lesson teaches the controls, the rules of the road and a test at the end. It works because a car does the same thing every time you turn the wheel. Learn the procedure once and it holds. Driving lessons fit cars. They are also what school has been giving children: set practices, passed on, checked. That method does not fit a horse.

A horse is different. It has habits. It shies. It is stronger on some days than you expect and no use at all on others. Nobody learns to ride from a manual. You learn by riding, and by being thrown, until the rider knows what the animal will do and what it will not.

A model is a horse. It gives every user the first, average answer. It sounds certain when it is wrong. It does not know your world. It solves the next step and not the one after. We call these habits Machine Defaults. Its abilities change in steps every few months, and nobody can give you a map of where it is good.

Not having the map has a cost. Dell'Acqua and colleagues at Harvard Business School gave 758 consultants at the Boston Consulting Group tasks to do with and without AI. On tasks inside the machine's ability, those using it finished 12% more tasks, 25% faster, at significantly higher quality. On a task chosen to sit just outside it, they were about 19 percentage points less likely to get the answer right. The edge was not marked. The authors call it the jagged frontier.

Wharton's Shaw and Nave found the same thing from the other side. In a preprint with over 1,300 participants and nearly 10,000 trials, a chatbot was right about half the time and confidently wrong the rest. When it was right, accuracy rose 25 points above

working alone. When it was wrong, accuracy fell 15 points below working alone. Confidence rose either way. They call it cognitive surrender.

	CAR	HORSE
THE MACHINE	Does the same thing every time	Has habits. Shies. Stronger some days than you expect
WHAT YOU LEARN	The controls and the rules of the road	A feel for this animal
HOW	Lessons, then a test	Hours in the saddle, and falls
WHAT CHANGES	Nothing under you	It gets stronger, and so must you

FIG. 6 · THE CAR AND THE HORSE

A child who rides directs the machine. A child who is carried takes what it gives.

Generating is asking the machine for an output and taking it. Creating is deciding what to make and for whom, directing the machine toward it, changing or throwing out what comes back, and checking the result against something real. A child can do either with the same tool in the same minute.

This is an invented example, so it shows the idea and not what happens. Two children ask a machine to make a maze game harder. The machine answers both the same way: "Done! I added more walls and a 30-second timer. Your game is now much harder." The first child says "ok thanks" and has the same maze thirty other children got. The second asks, "What else could 'harder' mean? Give me three different kinds." The machine offers walls that move, a treasure that runs away, fog that hides the path. The child thinks about her friends and says, "If the treasure runs every time, my friends will quit. Make it run only twice." She ends up with a game nobody else in the room made.

The second child used two of what we call the Four Moves. Generating Alternatives is asking for three answers and choosing, or throwing all three out. Tracing Consequences is asking, if this, then what? The other two are Revising Beliefs and Connecting Patterns. The machine did not have to be smarter for her to get that game. She had to stay the author. No adult appears in the example. The trial puts one in the room.

THE PASSENGER

Make my maze game harder.

Done! I added more walls and a 30-second timer. Your game is now much harder.

ok thanks

The same maze thirty other children got.

THE RIDER

Make my maze game harder.

Done! I added more walls and a 30-second timer. Your game is now much harder.

GENERATING ALTERNATIVES

What else could "harder" mean? Give me three different kinds.

1. Walls that move. 2. A treasure that runs away. 3. Fog that hides the path.

TRACING CONSEQUENCES

If the treasure runs every time, my friends will quit. Make it run only twice.

A game nobody else in the room made.

FIG. 7 · TWO CHILDREN, ONE MACHINE. AN EXAMPLE, NOT A TRANSCRIPT

The evidence that this difference matters is early, and it is thin, and it is about the tool, not the child. In one Turkish school, almost a thousand high-school students practised maths with GPT-4 (Bastani and colleagues, *PNAS*, 2025). Those given a version that simply answered scored 48% higher on practice problems, and 17% lower on the exam once it was taken away. Those given a version built to give hints scored 127% higher on practice and no different on the exam. The hints removed the loss. They did not add learning, even though the Harvard and Nigeria tutors did. It does not show that a child who directs the tool keeps more. That is a position, and the trial below is the way to test it.

What a child needs from the horse is a feel for it: what it does, where it shies, how far it can be trusted on a given day. A feel does not come from a lesson. It comes from riding, and from riding something that matters, because only then does it cost the rider when the horse goes wrong.

Creating brings both halves back

A child who creates meets both halves in the same afternoon.

Take the second child's maze. To make the treasure run twice and not every time, she had to count: how many times, how fast, how many moves before the timer runs out. That is half one, put to use. To decide that twice was right, she had to know how much her friends would put up with before they quit. That is half two: patience, trust, what makes someone keep playing. Nobody taught her either.

A child who makes something is pulled into everything around the making. She has to ask who it is for, what they need, what they will put up with or pay for, what the rules are, and what happens when it breaks. These questions do not arrive as a subject. They arrive as obstacles, one at a time, each with a cost if she gets it wrong. This is the day-to-day action that children once met in the family's work, with a machine in it to help.

The taxonomy that schools use to describe their own work, in the 2001 revision by Anderson and Krathwohl, puts creating at the top, above remembering, understanding, applying, analysing and evaluating. A school day built as a production line spends its hours on the lower rungs, because those are the rungs a line can carry. Creating sits at the top of the school's own list, and I take it as the ultimate form of learning. The proposal is to teach both halves this way, and to use AI to make that possible for every child, not for a few.

Two things look important, and the trial proposed below is built to test them: that what the child makes is real and gets used by someone, and that an adult is in the room, neither lecturing nor absent. The evidence that follows is why.



FIG. 8 · ONE AFTERNOON, TWO HALVES

The case for making is mixed, and the mix matters. A review of 30 studies of project-based learning, covering about 12,500 students in nine countries, found that it beat traditional teaching by 0.71 of a standard deviation on average (Chen and Yang, 2019). The best randomised trial we found, in five American school districts, raised the share of students passing Advanced Placement exams by a little under 8 points in the first year, among those who sat them. The second year's gain was larger, and the report calls it less conclusive (Saavedra and colleagues, 2021). The trial's funder promotes the method. And a review of 164 studies found that discovery with no help did worse than direct teaching, by 0.38 of a standard deviation, while guided discovery did better, by 0.30 (Alfieri and colleagues, 2011). So far, making wins when somebody scaffolds it.

Schools that teach enterprise by having children run mock businesses get mixed results. A Dutch mini-company programme, studied at one vocational college, lowered students' intention to start a business. A Swedish one was associated with a higher chance of starting a firm, up to sixteen years later. A Rwandan programme's gains had faded after three years. A randomised trial of a business game with 2,413 pupils aged 11 and 12 raised self-reported non-cognitive entrepreneurial skills and left knowledge unchanged. None of these measures whether a child understands money, audience or trust afterwards. They measure intentions, income or self-rated traits.

Some schools do teach pieces of half two as subjects, and it helps a little. A review of 76 randomised trials of financial education found gains of about 0.19 of a standard deviation in knowledge and 0.09 in behaviour. The behaviour gain was still there at two years, at 0.06, though the authors call the long-run evidence inconclusive (Kaiser and colleagues, 2022). A media-literacy course in Bihar, with over 13,500 students aged 13 to 18, raised the ability to spot false news by 0.32, and by 0.26 four months later (Amar and colleagues, 2026). The effects are small. A classroom can teach a piece of half two; it cannot put the child inside it.

As far as we found, nobody has tested what happens when children under eighteen make things for real customers with AI, with their understanding as the outcome. A Stanford review of AI in schools in 2026 read more than 800 papers and found about twenty with strong causal evidence. None was a study of American K-12 students.

Nobody has shown that creating with AI brings both halves back for children, which is why the proposal below is a trial. I hold that it does.

Creating is how a child gets a feel for the horse: when she has made something and it matters whether it works, it costs something when the machine goes wrong. So every child should get the chance.

Every child gets to ride

Every child needs to ride. Whether they do should not depend on who they are.


GIVE EVERY CHILD TOKENS, A LOT OF THEM, THROUGH AN ADULT INSTITUTION, AND WATCH WHAT THEY DO WITH THEM.

This is not about closing schools. Schools keep doing what they do: they mind children, keep them safe, and give them friends and a certificate. What is missing is a place where children create with the machine, with an adult in the room, because the adult is part of what makes it work. That can be a school, a library or a community centre.

The same tool, in the same minute, makes one child an author and another a passenger. If the strongest models, and the chance to create with them, reach some children and not others, the gap widens.

We can already see that exposure is uneven. The Direction Index, which we built, reads one provider's usage, from Anthropic's Economic Index, across 27 countries in one week of November 2025. It scores each country on two things: how heavily its working-age population uses the machine, measured against its size, and what the use is made of. That second part asks whether the person directs the machine or hands the job over, and whether the work is building or consuming. On the first, Israel is set at 100. The United States is 75, the United Arab Emirates 33, India 4.5 and Nigeria 4.1. On the second, India's use scores highest of the 27, at 68.6 against 56.0 in the United States, with about a twentieth of Israel's reach.

It is one provider's data, from users who skew technical, for adults. It says nothing about children. The data are Anthropic's, and we use Claude ourselves. It is still a picture of how uneven exposure is.



SOURCE: ILANO00 DIRECTION INDEX, FROM THE ANTHROPIC ECONOMIC INDEX, ONE WEEK IN NOVEMBER 2025. ONE PROVIDER, ADULTS ONLY. NOT A PICTURE OF CHILDREN.

FIG. 9 · REACH AND BUILDING IN 27 COUNTRIES, ONE PROVIDER

Handing a child a tool is not education. One Laptop per Child put computers into 318 rural Peruvian schools at random, with roughly one for every student. After fifteen months there was no gain in maths or language scores (Cristia and colleagues, 2017). In Lima, about a thousand laptops for use at home raised the children's skill with the laptop and had no effect on achievement (Beuermann and colleagues, 2015). In Romania, vouchers for home computers lowered grades by a quarter to a third of a standard deviation (Malamud and Pop-Eleches, 2011). The children got better at the machine and not at school. That is the reason to watch what children do with the tokens, and not only to hand them out.

Montessori's room was for the poorest children in Rome. The method spread as a name. By one 2022 census, about nine in a hundred of nearly sixteen thousand schools that use it are fully government funded. In the United States the revival began in 1958 with a school aimed at college-educated parents. The name was never protected, and a 1967 ruling found it generic. Nothing held the standard, and nothing held it open to every child.

A token is a piece of a word, and it is what the machine is paid in. Tokens are cheap, as the estimate earlier showed. The adults, devices and connectivity are not in that estimate, and nobody has priced them. Nobody does this today. Of the free education offers we found from the main providers, almost every one is for teachers, institutions or adults over eighteen. We found no government or charity scheme that gives children a token allowance, though our search was not exhaustive. A programme for children would be new.

It also has to run through an institution that carries the duty of care. Claude is for adults aged 18 and over, and an organisation that serves minors through its API has to verify ages, filter content, monitor use and tell users they are talking to an AI. Google's Gemini

is open from 13 with a personal or school account, and to younger children only through a parent's Family Link.

Measure how many tokens children use, how much of the use the child directed (whether the child decides, changes and checks, or hands the job over), and what was made and whether anyone used it. The Direction Index reads the second for countries. A Direction Report reads it for one child, marking the moments the machine pulled the work toward the average and what the child did about it.

The hard part is a fourth thing, and we do not have it. No instrument exists that tells whether a child understands money, audience or trust after making something. The trials above measured intentions, income and self-rated traits. Building that measure is part of the work.

To learn what happens, run it as a trial. One group of children gets tokens only, which is the call itself. A second gets tokens and guided creating, with an adult who is neither lecturer nor absent. A third gets neither. If the first group mostly hands the job over, that is a finding, and it shows what guidance adds. If some direct on their own, we want to know who and why.

	Tokens only	Tokens and guided creating	Neither
Tokens used	●	●	○
Share of use that is directed	●	●	●
What was made, and who used it	●	●	●
Does the child understand money, audience and trust?	○	○	○

● MEASURED ○ NOT APPLICABLE ○ INSTRUMENT STILL TO BE BUILT

FIG. 10 · THE TRIAL: THREE GROUPS, FOUR THINGS TO WATCH

Agree stop rules in advance, because unguided use can leave a child worse off, as the Turkish exam showed. Keep the data in aggregate, show each parent only their own child's record, sell nothing, and run it through an institution that verifies ages and carries the duty of care.

We ask governments and funders to pay for the trial and the adults, schools, libraries and community centres to host it, and AI providers to publish how their tools are used, by country and age band, in aggregate, so that nobody has to depend on one company's data, ours included.

What we can do ourselves today is less than this asks.

Careers


Nobody can list the jobs. The shape of the work is clearer than the titles.

Every parent asks what the child will do for a living. Any list of future jobs is a curriculum fixed in advance, which is what the production line did. The employers' own forecast shows why a list fails. The 2025 report from the forum at Davos expects 170 million jobs to be created by 2030 and 92 million displaced, a net gain of 78 million, alongside the 39% change in core skills. That counts formal jobs only, and it is an expectation, not a measurement.

What can be said is what the work will ask of a person.

It will ask her to direct the machine, with a feel for what it will and will not do, learned by riding. It will ask her to take an idea to a result that someone uses, without waiting to be handed a task. And it will ask her to know the human-made structure well enough to serve others inside it: what they need, what they will pay for, whom they trust.

Employers' own list of rising skills points the same way. Creative thinking, resilience and flexibility, curiosity and lifelong learning are all on it, beside AI and data skills.



NOT JOB TITLES

FIG. 11 · THE SHAPE OF THE WORK, THREE REINS

Footprint is the word for what you have made that others use. A child who has made ten things that real users used has a larger one than a child with ten good marks. That changes what a career is. A career used to mean selling hours. A person with a footprint is selling what she made, not only her hours.

In this paper a creator is anyone who puts a skill to use for others: a repair, a recipe, a lesson, a programme, a design. That can happen inside a company, beside one, or with none.

That is a long way from how work is organised today, and the numbers show how far. These counts are of narrower groups than the one this paper is about. Goldman Sachs counts about 50 million content creators worldwide and says only about 4% are professionals, earning more than \$100,000 a year. In the United States, 72.9 million adults did some independent work in 2025, by one industry survey, and over half of them only occasionally. India's NITI Aayog projects 23.5 million gig and platform workers by 2029 to 2030, up from 7.7 million in 2020 to 2021. In 2019 to 2020, more than half were drivers and sales staff, not makers. Work is moving from labourer to creator. For most it is not yet a living.

The proof of ability changes too. A degree tells you where someone studied, not what they can do. A record of real things made, and used by real users, tells you more. Some employers already say so: in India, 30% of employers named skills-based hiring, by dropping degree requirements, as a promising way to find talent, against 19% globally.

All of this is prediction. My answer to the parent's question, what should a child be doing now to be ready for work nobody can name, is creating: making things that others use, with a machine, and an adult beside her.

What we know, and what goes beyond it

Parts of this paper go further than the evidence. This section says which.

WHAT THE EVIDENCE SUPPORTS

- Children already use AI chatbots. In the United States, 64% of teens do. In Britain, 56% of children aged 8 to 17 have.
- The same tool can leave a learner worse off when it does the work. In a Turkish school, exam scores without the machine were 17% lower than the comparison group's. The version built to give hints removed the loss and added no learning.
- Adaptive AI tutoring raised learning in trials of students from secondary age up: 194 Harvard students, and over 600 secondary students in Nigeria, the second as a working paper.
- School administrators described school in production terms and applied efficiency methods to it: Cubberley in 1916, and the years of 1911 to 1925 that Callahan studied. The graded classroom itself dates from 1848.
- Taking part in a family business, or being exposed to inventors, raises a child's chances of entrepreneurship, invention and income.
- Project-based teaching beat traditional teaching on average, in studies some of which have interested funders. Unassisted discovery did worse than direct teaching, and guided discovery did better.
- Teaching pieces of half two in a classroom helps, modestly.
- Employers expect large change in skills, and a net gain in jobs.

WHAT GOES BEYOND THE EVIDENCE

These are my positions, and no study yet settles them. Half two is learned by taking part, as money, trust and rules, and not only in the narrow cases that have been measured (entrepreneurship, invention and income). Children who create real things with AI meet both halves and build a feel for the machine, though no study has looked at children under eighteen doing this. Whether a child generates or creates decides what she keeps, though the one study behind that compares answering with hinting, which tests the tool and not the child. Attention that adapts to one learner is cheap enough to remove the reason for the production line, though the trials are of older learners. Labour and capital are merging where work is procedure, judgment is what remains, and work moves toward smaller units. A record of real things made and used, a footprint, will count for more than a degree. Much of what we call intelligence is exposure, which goes further than the data.

WHAT WOULD SHOW WE ARE WRONG

- Children given guided creating understand money, audience and trust no better than a comparison group.
- The group given tokens only does as well as the guided group, so guidance adds nothing.
- Unguided use leaves children worse off on tasks without the machine, and guidance does not fix it.
- What children make turns out to depend on the machine and not on the child, so that the measure of direction tells us nothing.

WHAT WE DO NOT KNOW

- Whether anyone can measure what a child understands about money, audience and trust after making something.
- How many children in India use AI. We found no national figure.
- What the adults, devices and connectivity cost.
- Whether any effect lasts, or carries from one country to another.
- Whether the Direction Index, built on one provider's data for adults, tells us anything about children.

Every number has a source in the appendix, which also says which sources we read only as summaries.

What Ilanooo does today

Ilanooo is small. Two things exist today.

The Direction Index is public. It reads one provider's usage across 27 countries and scores each on how heavily its working-age population uses the machine, against its size, and what the use is made of. Appendix B sets out what it can and cannot say.

The Direction Report is a page about one child. The child builds a small game with an AI developer that pulls every game toward the average on purpose. The parent gets a report on the moments the child kept their own direction and the moments the machine took over. A parent can get one today.

Ilanooo is building programmes for children and students, and in time for schools. A workshop, a community and a trial do not exist yet. We want to run the first trial, through an institution that carries the duty of care. None has started. It needs funders, and it needs the adults.

The horse is already here. Whether every child learns to ride is a decision, and it has not been made.

Sources and notes

P means we read the primary text or full abstract. S means we read a summary only and have not read the primary text. E is our own estimate. O is our own data.

PART	CLAIM	SOURCE	HOW WE READ IT
Prologue	Room opened 6 Jan 1907 at 58 Via dei Marsi, about 50 children aged 3 to 6	Montessori, <i>The Montessori Method</i> (1912) ch. 2; Kramer, Maria Montessori (1976)	P
Prologue	Owners 'not prepared to spend a penny on toys or equipment'; assistant Candida Nuccitelli, about 40, porter's daughter, untrained	Kramer (1976); Montessori (1912)	P
Prologue	Writing about ten months after opening; reading within days	Kramer (1976)	P
Prologue	Four houses by the end of 1908; sponsor and doctor fell out over the publicity; porter told not to let her in	Montessori (1912) list of houses; Kramer (1976)	P
Prologue	State school at 58 Via dei Marsi today	Opera Nazionale Montessori listing	S
Prologue	About 15,763 schools in 154 countries, about 9% fully government funded	Global Montessori Census 2022 (Debs et al., <i>J Montessori Research</i> 8(2)); triangulated estimate, 42% response	P
1	Markets and bureaucracies already work by lossy summary; a price 'compresses vast swathes of detail into a simplified but usable representation'	Farrell, Gopnik, Shalizi & Evans, <i>Science</i> 387:1153–56 (2025); author's full-text version (no page numbers)	P
1	1299: the Florentine money-changers' guild (<i>Arte del Cambio</i> , article 102) told members to write sums in words, not abacus figures; reason unstated; not a general ban	Nothaft, 'Medieval Europe's Satanic Ciphers' (Oxford, full text); Camerani Marri ed. (1955)	P
1	Madhava and the Kerala school: infinite series for sine, cosine and arctangent by about 1400, over two centuries before Gregory, Newton and Leibniz; Yuktibhasa c. 1530 gives proofs	MacTutor biographies of Madhava and Jyeshthadeva	P
2	Learning by observing and pitching in, described in Indigenous-heritage communities of the Americas	Rogoff, <i>Human Development</i> 57(2-3) (2014); concept read at the LOPI site	P

PART	CLAIM	SOURCE	HOW WE READ IT
2	'Only in the single US middle-class community of Orchard Town were children not routinely engaged in chores'	Lancy, The Anthropology of Childhood (CUP), chapter overview; page not found	P
2	English apprenticeship contracts: 'The details of what was to be learned were not specified'; 'keep his master's secrets'; about one teenage boy in ten apprenticed in 1600 and 1700	Wallis (2019), Apprenticeship in England	P
2	Being raised by entrepreneurs about twice the effect of being born to them (adoption data)	Lindquist, Sol & Van Praag (2015), J Labor Econ 33(2)	P
2	A relative's business helps only if you worked in it (sales about 40% higher); the bare family tie was small and insignificant	Fairlie & Robb (2007), J Labor Econ 25(2) (2004 IZA working paper read; check published tables)	P
2	Over 1.2 million inventors; children who moved young to places with more inventors became inventors at higher rates	Bell, Chetty, Jaravel, Petkova & Van Reenen (2019), QJE 134(2)	P
3	Quincy Grammar School, Boston, 1848, principal John D. Philbrick: usually credited as the first American school with one grade, one room, one teacher; twelve classrooms of about 56 pupils	Encyclopedia.com (Gale, American Eras); Tyack's text not read	S
3	Global adult literacy about 12% in 1820 (1 in 10 over age 15), about 87% now	Our World in Data, Literacy (van Zanden et al. 2014; UNESCO)	P
3	Callahan: 'efficiency' procedures applied to schools 1911 to 1925	Callahan, Education and the Cult of Efficiency (1962), ch. 5	P
3	Cubberley: 'Our schools are, in a sense, factories ...' (p. 338)	Cubberley, Public School Administration (1916), archive.org scan	P
3	1912 critic: 'terms that prevail in the factory and the department store'	B. C. Gruenberg, American Teacher (Sept 1912), quoted in Callahan	P
3	Tyack: the graded, centralised, professionally run school was built in the late 19th and early 20th centuries as 'the one best system' for the industrial city	Tyack, The One Best System (1974), via History of Education Quarterly review essay; book not read	S

PART	CLAIM	SOURCE	HOW WE READ IT
4	39% of core skills to change by 2030 (down from 44% in 2023); analytical thinking top skill, called essential by seven in ten companies; survey of 1,000+ employers covering 14+ million workers; clerical and secretarial largest absolute decline; 170m jobs created, 92m displaced	World Economic Forum, Future of Jobs Report 2025 (PDF read in full)	P
4	Typists about -34%, data entry keyers about -25%, cashiers about -200,600 (2025-35); 'automated systems, including AI ...' (2024-34 release)	US Bureau of Labor Statistics projections tables and release	P
4	Harvard physics: N = 194, gains more than double, median 49 vs 60 minutes	Kestin et al. (2025), Scientific Reports 15:17458	P
4	Nigeria: over 600 students, six weeks, +0.31 SD; working paper	De Simone et al. (2025), World Bank Policy Research Working Paper 11125	P
4	Tokens for a child using ~100k tokens a day: about \$0.47 to \$13.20 a month	Our estimate from current official list prices; assumptions: 70/30 input/output, 30 days, no caching	E
5	UNESCO AI competency framework for students (8 Aug 2024): four dimensions, 12 competencies. UAE: AI a formal public-school subject, kindergarten to Grade 12, from 2025-26	UNESCO (official page); Gulf News (the official u.ae page returned 404)	S
5	64% of US teens use an AI chatbot; about 3 in 10 daily; 54% for schoolwork	Pew Research Center (Dec 2025, Feb 2026), n = 1,458	P
5	56% of children 8 to 17 have used AI; 66% of 16 and 17 year olds	Ofcom, Children and Parents: Media Use and Attitudes (21 May 2026), n = 3,426	P
5	Princeton faculty vote, 11 May 2026, one dissent; proctoring banned 133 years; 29.9% cheated, 0.4% reported	The Daily Princetonian (May 2026)	P
5	758 consultants; inside the frontier 12.2% more tasks, 25.1% faster, quality significantly higher; outside, about 19 percentage points less likely correct (the published abstract says '19% less likely'; the body and working paper say percentage points)	Dell'Acqua et al., HBS WP 24-013; Organization Science 37(2) (2026)	P
5	1,300+ participants, ~10,000 trials; +25 points when AI right, -15 when wrong; confidence up	Shaw & Nave, 'Thinking, Fast, Slow, and Artificial' (preprint, SSRN)	S

PART	CLAIM	SOURCE	HOW WE READ IT
5	Turkish school: GPT Base +48% practice, -17% exam; GPT Tutor +127% practice, no significant exam change	Bastani et al., PNAS 122(26) (2025); correction was affiliation-only	P
6	Project-based learning: 30 articles, 46 effect sizes, 12,585 students, 189 schools, 9 countries, $d + 0.71$	Chen & Yang (2019), Educational Research Review 26 (abstract read; full text blocked)	P
6	Five US districts: year one +7.6 points in AP passing among exam-takers (+4 across all students); year two about +10 for second-year teachers' students, called 'less conclusive', attrition above WWC thresholds; funded by Lucas Education Research, which also funded the curriculum	Saavedra et al. (2021), Knowledge in Action, full report	P
6	164 studies: unassisted discovery $d = -0.38$; enhanced discovery $d = +0.30$	Alfieri et al. (2011), J Educational Psychology 103 (abstract)	P
6	Dutch mini-company: intention negative, one vocational college, instrumental variables (2010). Swedish: associated with more firm starts, up to 16 years, matched design (2015). Rwanda: +6 points in entrepreneurship at one year, faded by year three, employment 6 points lower; working paper (2023). BizWorld: 2,751 recruited, 2,413 analysed; knowledge unchanged, self-reported non-cognitive skills up (2014)	Oosterbeek et al. (2010), Eur Econ Rev 54(3); Elert et al. (2015), JEBO 111; Blimpo & Pugatch (2023), EdWorkingPaper 23-861 / IZA DP 16489; Rosendahl Huber et al. (2014), Eur Econ Rev 72	P
6	Financial education, 76 RCTs: +0.186 SD knowledge, +0.094 behaviour; 0.057 SD at two years or more (7 RCTs, interval excludes zero); authors call long-run evidence inconclusive	Kaiser, Lusardi, Menkhoff & Urban (2022), J Financial Economics 145(2) (NBER working-paper version read)	P
6	Bihar: 583 villages, over 13,500 students aged 13 to 18; +0.32 SD discernment; +0.26 SD at four months (a subsample of 2,059)	Amar, Badrinathan, Chauchard & Sichart, American Political Science Review 120(2), May 2026	P
6	800+ papers, about 20 with strong causal evidence, none a US K-12 study	Stanford SCALE, Evidence Base on AI in K-12 (2026)	P
6	Revised taxonomy has six levels, Create highest; the authors say the strict hierarchy has been relaxed	Anderson & Krathwohl (2001), via Krathwohl (2002), Theory Into Practice	P
7	Direction Index: 27 countries, 623,460 conversations, 13 to 20 November 2025; reach and build scores	Ilano00 Direction Index, from the Anthropic Economic Index, fourth report, 'Economic Primitives' (15 Jan 2026)	O

PART	CLAIM	SOURCE	HOW WE READ IT
7	Peru: 318 schools, no gain in maths or language; Lima: no achievement effect; Romania: grades down a quarter to a third of an SD	Cristia et al. (2017); Beuermann et al. (2015); Malamud & Pop-Eleches (2011)	P
7	Whitby School 1958 for college-educated parents; 1967 ruling that 'Montessori' is generic	Whitescarver & Cossentino (2008); US Trademark Trial and Appeal Board (1967)	S
7	Claude 18+; API organisations serving minors must verify age, filter, monitor, disclose AI. Gemini 13+ with a personal or school account	Anthropic Help Center (16 Mar 2026); Google Gemini help pages	P
7	No government or charity token scheme for children found; free offers mostly for teachers, institutions or adults 18+	Official provider education pages; search not exhaustive	S
8	About 50m creators; 'only about 4% ... are deemed professionals, meaning they pull in more than \$100,000 a year'	Goldman Sachs Global Investment Research (19 Apr 2023)	P
8	72.9m US adults did independent work in 2025 (27.6m full-time, 7.9m part-time, 37.4m occasional); survey of 6,474 adults, April 2025	MBO Partners, State of Independence 2025	P
8	India gig and platform workers 7.7m (2020-21), projected 23.5m (2029-30); drivers and sales staff more than 52% in 2019-20 (base 6.8m)	NITI Aayog (June 2022), p. 22	P
8	India: 30% of employers named skills-based hiring by dropping degree requirements as promising for finding talent, against 19% globally	World Economic Forum, Future of Jobs 2025, pp. 53 and 68	P

How the Direction Index is built

The Direction Index is Ilanooo's own measure. It is computed from the Anthropic Economic Index open dataset, the fourth report, "Economic Primitives", published 15 January 2026. The data window is 13 to 20 November 2025. The Index covers 27 countries and 623,460 conversations. It has four parts, as the Index page states them.

Penetration. Anthropic's AI Usage Index (AUI) is a country's share of Claude usage divided by its share of the working-age population. We normalise it to a 0 to 100 scale against the highest national value, Israel, at 4.90, which becomes 100.

Build and direct score. This is the Index's directing-quality signal, built from two measured parts weighted equally. The first is directing mode: whether usage is collaborative directing or a delegated hand-off, from the collaboration facet of the data. The second is task mix: whether the work is building or consuming, from the task facet.

The direction score. The Index page gives it as penetration multiplied by the composition term, on a 0 to 100 scale. Multiplying forbids a large volume of shallow use from standing in for substance. Part 7 does not rely on the combined score. It uses penetration and the build and direct score as published, and we have not recomputed the combined score from the columns shown.

The bifurcation index. Build quality divided by penetration, with a small constant to keep the curve bounded near zero. It asks whether directing is spread through a population or held by a thin group.

COUNTRY	REACH (ISRAEL = 100)	BUILD AND DIRECT SCORE	CONVERSATIONS SAMPLED
Israel	100.0	54.6	7,920
United States	75.3	56.0	219,666
United Arab Emirates	32.9	60.9	3,883
India	4.5	68.6	58,098
Nigeria	4.1	58.9	7,629

LIMITS.

- One provider. Claude's users skew toward technical and professional work. ChatGPT, far larger and more consumer-tilted, does not publish per-country composition.
- Adults, and one week. The Index says nothing about children.
- Small samples. New Zealand, Norway and Ireland are flagged, each with fewer than 3,000 conversations.

- The Index measures the directing of a country's serious AI-working population. It is not a measure of whole-population adoption.
- We use Claude ourselves.