

311 INSTITUTE · INTELLIGENCE UNIT

FUTURE OF WORK: ONE PERSON, 10,000 HANDS

How AI agents and robots multiply
what one person can produce, Rome to
2100

REPORT

Deep research report

BY

Matthew Griffin · London

DATE

13 October 2026

ABOUT THE 311 INSTITUTE

The 311 Institute, a London-based global leader in Futures, Deep Futures and Dynastic Futures advisory, exists to serve and help leaders anticipate, explore and shape every future.

History has repeatedly shown that organisations are made or broken by disruption — made by the shifts they saw and acted on in time, and broken by those they missed, misunderstood or met too late.

The VUCA-MAX world runs on accelerating, complex change — hard to anticipate, harder to answer. Almost half the world's CEOs believe their companies will not be economically viable within a decade without reinvention. That widening gap — seeing the future, understanding it, acting in time — is the one we exist to close.

With a track record of trust and excellence built advising the world's largest organisations, we close that gap in a way no one else can. Where others offer opinion we offer a record, where others see slices of the future we see the whole horizon, and where others offer partial solutions we offer everything.

Used by the world's most consequential organisations — among them Walmart, the world's largest employer, alongside ABB, Accenture, Adidas, Allianz, AON, ARM, AT&T, Auchan, BCG, BlackRock, Citigroup, Coca-Cola, DP World, Emirates Airlines, EY, the FAA, Google, HP, Huawei, KPMG, LGT, Nokia, Novartis, PepsiCo, PwC, Roche, RWE, Samsung, Siemens AG, Standard Chartered, T-Mobile, UBS, Visa, Workday, Worldpay and the United Nations, together with royal households and G7, G20 and G77 leaders. We advise boards, governments and investors across every industry and at every scale, helping countries grow and enterprises become Centennial companies built to endure beyond a hundred years.

In a world that punishes the unprepared, the gap between seeing the future and shaping it is all too often the one that determines who thrives and survives, and who doesn't. We close that gap.

ABOUT THE AUTHOR

**Matthew Griffin**

Founder & Futurist in Chief · 311 Institute, London

Matthew Griffin is a multi-award-winning futurist and expert in disruption and innovation, geopolitics, leadership and technology, described as a "walking encyclopaedia of the future" and a "futurist polymath." A 15-time best-selling author of the "Codex of the Future" series, he is the Founder and Futurist in Chief of the 311 Institute — a global Futures and Deep Futures advisory firm working with royal households, world leaders, G7, G20 and G77 governments, NGOs and multinational mid- and mega-cap firms to help them explore, shape and lead the next 50 years of business and society.

An award-winning YouTube creator with over a million followers and unrivalled global reach, Matthew is a highly sought-after international keynote speaker, lecturer and mentor. He collaborates with global leaders through the United Nations Alliance of Civilizations (UNAOC) and the United Nations General Assembly (UNGA) to shape initiatives such as the UN's AI for Humanity programme, the UN Conference of the Parties (COP) and the World Economic Forum in Davos.

As the former Global Head of Cloud, National Security and Enterprise Sales for companies including Atos, Dell-EMC and IBM, Matthew has a proven track record of building multi-billion-dollar business units and turning failing divisions into market leaders. His ability to identify, analyse and communicate the implications of hundreds of emerging technologies and trends is unparalleled, and his insights are trusted by many of the world's most respected organisations.

Regularly featured in the global media — including the AP, BBC, Bloomberg, CNBC, Discovery, Forbes, Khaleej Times, the Telegraph, TIME, ViacomCBS, WIRED and the WSJ — Matthew's mission is to help organisations create a fair and sustainable future whose benefits are shared by everyone, irrespective of ability, background or circumstance.

CONTENTS

PART I

What this report is for 09

PART II

How we built the numbers 17

PART III

One body, one output 28

PART IV

The Roman family of four 34

PART V

Mills, monks and the long flat line 40

PART VI

Ships, coal and steam 47

PART VII

The century of energy servants 53

PART VIII

The average global family, 2000 to 2026 59

PART IX

The command equation 65

PART X

Swarms of agents 72

PART XI

Fleets of robots 80

PART XII

Infinite worlds 87

PART XIII

What binds 94

PART XIV

One person in 2100 101

PART XV

Where the wealth accrues 113

PART XVI

Work, earnings, tax and economies	122
-----------------------------------	-----

PART XVII

Moves for every stakeholder	129
-----------------------------	-----

PART XVIII

The output data book	136
----------------------	-----

CLOSING REMARKS

Strategic summary	144
Questions to take away	145

ANNEX

Notes and method	148
Assumptions and caveats	149
Source limitations	150
Glossary	151
Abbreviations	159
Source ledger	165
Terms & conditions	180

THE 311 INSTITUTE READ

IN BRIEF

This report asks how much one person can make. It follows the answer from a Roman farm family in 1 AD to an ordinary person in 2100. It counts work in one unit, the Human Labour Equivalent (HLE), which is one adult's working year. It counts money separately, in dollars. It splits output three ways. Digital output is software, data, services and content on ordinary screens, including AI-made music and video. Physical output is every physical product or service, from medicines to rockets. Virtual output is anything made for VR, AR and spatial computing.

THE OBVIOUS READ

The usual view is that AI and robots are just faster tools. People keep the same jobs, some jobs go, and output per hour rises a little each year, as it did with computers. Pay stays hours sold times a price.

THE FULLER PICTURE

Our view is that the unit itself breaks. One person stops selling one person's work and starts directing the work of many, as only slave owners, lords and mill owners could before. By 2100 one skilled person could direct about 260 HLE a year in a setback, 700 in a long drift, 30,000 if AI agents spread widely, or 2 million in an explosive window. Work becomes plentiful, its price falls, and wealth goes to whoever owns what stays scarce.

EXECUTIVE SUMMARY

This report is for leaders, investors, teachers and curious readers who want to know how much one person will be able to make, and who will keep the gains. For about 300,000 years one human body made about one human's worth of work. That is about 75 watts, or one HLE a year. Every gain since has come from directing more than one body. A typical Roman farm family of four did about 2.5 HLE of work a year, and the average family, counting the rich, about 10. One of Rome's richest freedmen, Caecilius Isidorus, directed about 33,000 HLE through 4,116 slaves and 3,600 pairs of oxen. For 2,000 years the multipliers (slaves, mills, ships, coal and engines) were owned by a few. The European family's useful work reached about 23 HLE by 1800, 100 by 1900 and 290 by 2000, almost all of it from fossil-fuelled machines. Today the average person on Earth gets about 38 HLE of work a year.

The new step is delegated thinking (AI agents that do mental work for us) and dexterous robots (machines with hands that can do varied physical work). This report's Command Equation says a person's output is their own work plus all the work they direct. That is units times hours times speed times the share of work accepted. Four limits cap it: how much a person can direct, their budget, energy and what buyers want. A heavy user paying \$180 a month already directs about 18 HLE of AI agent work, the work of about 18 full-time people, for about \$270 of electricity. For heavy users, attention is the limit. For most people, money is. Up to 2050 we give odds. By 2050 a skilled person directs about 120 to 6,100 HLE a year, with about 660 in the central case. After 2050 we give branches without odds. In 2100 one skilled person could direct about 260 HLE a year after a setback, 700 in a long drift, 30,000 if agents spread widely, or 2 million in an explosive window.

If agents spread widely, what that skilled person directs in 2100 is about 74% digital, 25% virtual (VR, AR and spatial worlds) and about 1% physical. Virtual work is only about 1% of AI work today. Output per person rises about a hundred times by 2100 if agents spread widely, but measured income rises only about fourteen times. Plentiful work makes prices fall. History says the owners of a new multiplier gain first. The share of income paid to workers is already falling, and the first effect on jobs is fewer junior hires. The moves that pay in every branch are to measure delegated work, train people to direct and check it, protect junior jobs, own something scarce and spread ownership of the multipliers.

01

PART ONE

WHAT THIS REPORT IS FOR

What this report is, who it is for, and how to read it.

01

How much can one person produce, and why does that question matter now?

For all of history, one person has made about one person's worth of work. AI agents and robots are about to break that rule. This report measures what one person can make and direct, from Rome to 2100.

SCOPE

Rome to 2100

One person

Not welfare or happiness

This is a deep research report from the 311 Institute. It asks one simple question: how much can one person produce? We follow the answer for more than two thousand years. We start with a Roman farm family in 1 AD and end with an ordinary person in 2100. Along the way we look at oxen and slaves, mills and steam, computers and the internet. Then we look ahead to AI agents and robots. At each step we ask three things, starting with how much work got done. Who or what did it, a person or their helpers? And who kept the money it earned? The answers change a great deal over time, and they are about to change again.

Span

1 AD–2100

Unit

HLE

Branches after 2050

4

WHAT THIS REPORT IS

The report is told as a story, but it is built on numbers. Each number has a source, and each source has a grade. Where we are making an estimate about the future, we say so. Where experts disagree, we show both sides and explain which one we use.

WHY IT EXISTS NOW

For almost all of history one rule held: one person made about one person's worth of work. A farmer could dig only so much soil in a day, and a clerk could copy only so many letters. Pay, taxes and pensions are all built on that rule. You sell your hours, and the hours are yours to sell.

That rule is now starting to break. An AI agent (software that can carry out a task on its own, step by step) can write code or draft a report. It does this while its owner does something else. A heavy user can already run ten of them at once. Soon, robots may do the same for physical work. So the useful question is no longer only what a person can make with their own hands, but also what they can direct.

Nobody measures that yet. Official figures count hours worked and money earned, but they do not count the work a person sets in motion through machines and software. This report tries to fill that gap, so it builds one yardstick that works for a Roman ox, a steam mill and an AI agent.

WHO SHOULD READ IT

We wrote it for leaders in government and business, who set the rules and make the plans. We wrote it for investors, who need to know where the gains will collect. We wrote it for teachers and students, who are preparing for jobs that may look very different. And we wrote it for curious readers who simply want to know what is coming.

You do not need to know any economics to read it. Every technical term is explained in plain words the first time it appears. The annex at the back has a glossary of terms and a list of abbreviations.

THE QUESTIONS IT ANSWERS

The report answers five main questions. First, how much could one person produce in the past, and what made that number grow? Who owned the things that multiplied a person's output? How much can one person direct today? How much might they direct in 2050 and in 2100? And who will keep the gains: the person, their employer, the owners of the machines, or the state?

HOW THE REPORT IS ORGANISED

The history chapters come first. They run from the unaided human through the Roman family, mills, ships, coal and steam, to the average family of 2026. They show how each new helper raised output, and who got the gains. Often it was not the people doing the work.

The future chapters come next, and they set out the Command Equation, a simple formula for what one person can direct. Then they look at swarms of AI agents, fleets of robots and virtual worlds in VR and AR. They ask what limits growth, and what one person might direct in 2100.

The consequences chapters ask what all this means for real lives. They cover where wealth collects, what happens to work and tax, and what governments, firms, investors and people can do now. The data book at the end of the report holds the full tables, decade by decade. The annex holds a note on every source, the glossary and the abbreviations.

READING PREDICTIONS, FIGURES AND SOURCES

Three features appear again and again through the report: prediction blocks, figures and source markers. Here is how to read each one.

Prediction blocks. Each one makes a clear claim about the future that can be checked. Each block shows a figure labelled CONFIDENCE, such as 40%. That is the chance we give the call: our honest estimate that the claim comes true. A short basis says why, and names an early sign to watch for. The block ends with a tripwire that starts "Proved wrong if". The tripwire tests the same thing the claim says, so you know exactly what would show the call was wrong. Every prediction is settled by 2050 at the latest.

Figures are the charts and tables in the report, and each one has a title, a source line and labels. Many charts use a log scale, where each step up the side is ten times the step below. That lets a Roman farm and a robot fleet in 2100 fit on the same page.

Source markers are small links in the text, and each one points to a note in the annex that says where a number came from. Each note also carries a trust tier. Primary means the number comes straight from the body that measured it. Corroborated means two or more good sources agree on it. Single-source means only one source has it, so treat it with care. Unverified means the figure is disputed or self-reported and has not been checked. Foresight means it is our own estimate about the future, not a measured fact. In tables, our own calculations are marked "our estimate".

KEY IDEAS IN PLAIN WORDS

The report's unit is the HLE, short for Human Labour Equivalent. It is one adult's working year, counted in one of two ways. The physical way is about 75 watts of useful work for 2,500 hours, roughly what a healthy adult can keep up^[8]. The time way is 2,500 accepted hours of work, done by a person, an AI agent or a robot. Money is never counted in HLE, so when we say what something is worth, we use 2025 US dollars.

Three words sound alike but mean quite different things. Output is how much work gets done, and productivity is output per hour worked. Capture is who keeps the money that the output earns. A person can direct far more output while their pay stays flat, if someone else captures the gain. History shows this happening again and again, from Roman estates to steam mills.

Output comes in three kinds. Digital output is everything digital that one person creates. That means software, apps, data, analysis and digital services. It also means content on ordinary screens, such as text, images, music, video and games. So a song or an article made by AI counts as digital. Physical output is every physical product or physical service. It runs from medicines to rockets, and from food and furniture to buildings and moved goods. Virtual output is anything made in or for virtual reality (VR), augmented reality (AR), mixed reality and spatial computing. VR puts you inside a computer-made world through a headset. AR lays computer images over the real world, often through smart glasses. Spatial computing means apps and worlds you move around in, not just look at on a flat screen. Virtual output is tiny today, about 1% of AI work.

Two new kinds of helper matter most in this report. An AI agent takes a goal and works through the steps alone. A humanoid robot is a robot shaped roughly like a person, so it can use tools and spaces built for people. The Command Equation puts these together and adds up a person's own work and everything they direct. Then it caps the total at the first limit that bites: their attention, their budget, their energy or the demand for what they make.

Finally, the report treats the future in two parts. Up to 2050 we show a band of outcomes. Its low edge is the long drift path, its middle is the agentic diffusion path, and its high edge is the explosive window path. For this period we give probabilities, but after 2050 we show four branches with no odds: setback, long drift, agentic diffusion and explosive window. Nobody can honestly put odds on 2100, so we show what decides the branch, and what to do in every case.

THE HEADLINE ANSWER

So how much could one skilled person direct, if they spent a tenth of an average income on AI and robot work? In 2026 the answer is about 56 HLE a year. In 2100 the answer depends on the

branch. One skilled person could direct about 260 HLE a year in the setback branch or 700 in long drift. That rises to 30,000 in agentic diffusion and 2 million in the explosive window ^[93].

BRANCH	WHAT HAPPENS	HLE A YEAR, 2100
Setback	A hard shock in the 2040s costs a decade, then slow recovery	260
Long drift	AI and robots spread slowly, growth stays modest	700
Agentic diffusion	Agents spread widely and get cheap, people learn to direct them	30,000
Explosive window	AI does most tasks and growth runs very fast for a time	2 million

Source: 311 Institute analysis 2026 (foresight). One skilled person spending a tenth of an average income.

WHY 56 HERE BUT 18 LATER

The 2026 figure of about 56 HLE is for an ordinary skilled person, not a heavy AI user. They spend a tenth of an average income, about \$2,600 a year. Four-fifths of it, about \$2,080, buys AI work at about \$120 for each HLE. That is about 17.6 HLE. The last fifth, about \$520, buys robot work. Robot work still costs about \$15,000 for each HLE, so it buys almost nothing yet, about 0.03 HLE. Add their own work, which is 1 HLE. Then add their share of everyday machines, such as motors, power stations and factories, which is about 37.2 HLE. So 1 plus 17.6 plus 37.2 gives about 56. For this person money runs out first, so the budget is the limit.

The next chapter, on method, works through a different case. A heavy user on a flat \$180-a-month plan runs ten agents and directs about 18 HLE of accepted AI work. That is the work of about 18 full-time people. For them money is not the limit, because the plan is a flat fee. Attention is the limit, since one person can only set and check so much work.

So compare like with like. The ordinary person gets about 17.6 HLE of AI work, close to the heavy user's 18. The heavy user spends all of their money on AI work, so they get a little more for less. The 56 is bigger because it also counts their own work and their machines. The two numbers answer different questions, and both are right.

The gap between the 2100 numbers is huge, but that width is not a failure of the method. It is the honest finding. The result depends on choices and shocks that have not happened yet. The rest of this report shows where each number comes from, and what would push the world towards one branch or another.

CHAPTER SUMMARY

This report measures how much one person can make and direct, from a Roman farm to 2100. It counts work in HLE and money in dollars. It keeps output, productivity and capture apart. Up to 2050 it gives a percentage chance for each call. After 2050 it gives four branches with no odds, and the signs that show which one we are on.

The old rule was one person, one person's output. This report asks what will take its place.

CHAPTER CLOSE

WHAT NEXT?

What this report is for

QUESTIONS TO ASK

- 1 How much work could you direct today if you tried?
- 2 Which of the four branches would your plans survive?



PART TWO

HOW WE BUILT THE NUMBERS

The unit, the data, the formula and the rules behind every figure.

02

How can one yardstick measure a Roman ox, a steam mill and an AI agent?

Every number in this report comes from a source, a stated method or a clear assumption. This chapter shows each step, so a reader can check our work or disagree with it.

SCOPE

Data

Formula

Branches

Money is a poor ruler for work across two thousand years. Prices change and currencies vanish. An hour of light from a bulb today costs a tiny part of what an hour of candlelight cost in 1800. So money hides how much more light we now make. Money also mixes up output with capture, because if a firm keeps more of the gain, workers' pay can stay flat even when output soars. We need a ruler that means the same thing in a Roman field, a Victorian mill and a data centre. This chapter builds that ruler. Then it shows the data, the formula and the rules behind every number, so you can check each step.

Unit

HLE

Steps

100, then 10 years

Sources checked

Oct 2026

COUNTING WORK, NOT MONEY

So the report uses the HLE, a unit of work rather than money. The HLE can be counted in two ways. The physical way is about 75 watts kept up for 2,500 hours. That is 190 kWh, or 0.684 GJ (gigajoules), about the energy in 20 litres of petrol ^[8]. We use it for muscles, animals, mills and motors. The time way is 2,500 accepted hours of work by a person, an AI agent or a robot. We use it for tasks, where watts tell you little.

Money still matters, so we report it next to the work. It is always shown in 2025 US dollars, and never turned into HLE ^[9]. Keeping the two apart is the whole point, because it lets us show when output grows faster than income, and who keeps the difference.

A ROMAN FAMILY AS YARDSTICK

Our baseline is a Roman family of four in 1 AD: two adults and two children on a small farm. We use a family, not a lone worker, because for most of history work was done by households. Children helped in the fields, and old people spun and cooked. A family also makes the numbers easy to picture.

An average Roman family had an income of about \$9,900 a year in 2025 dollars. Its useful work, counting people, animals and mills, came to about 10 HLE a year. A poor smallholder family made only about 2.5 HLE of work, worth roughly \$2,500.

From 1 AD to 2000 we follow Western Europe as the spine of the story. The data are richest there, and Rome is the starting point. We note other places, such as Song China, where they led. From 2000 we switch to an average global family, using world averages. The big changes of this century are global, so a single country would mislead.

CENTURIES FIRST, THEN DECADES

We move in 100-year steps from 1 AD to 2000. Older data are too rough for finer steps, and change was slow. From 2000 we move in 10-year steps to 2100. Change is now fast, and the data are much better. The data book at the end of the report shows every step in full.

SOURCES FOR THE HISTORY NUMBERS

Income per head before 2000 comes mainly from the Maddison Project, the standard long-run record of world incomes ^[1]. For Rome we also use the estimates of Lo Cascio and Malanima ^[2]. For

power and energy we use Vaclav Smil ^[8] and the Energy Institute ^[27]. From 2000 we use the World Bank ^[26], and for population the United Nations ^[86].

These sources use old units, so we convert them. Maddison counts in 1990 Geary-Khamis dollars, a unit built to compare what money could buy in different countries. We multiply those by about 2.25 to reach 2025 dollars. The World Bank uses 2021 PPP dollars (dollars adjusted so they buy the same amount in every country). We multiply those by about 1.17 to reach 2025 dollars.

The error bars are wide, and we say so. Before 1500, good estimates often differ by about twofold. For Roman Italy, Maddison gives about 809 dollars a head in 1990 money. Lo Cascio and Malanima give a higher range of 1,300–1,800 ^[2]. Our series uses about 1,100, between the two estimates. Treat any number before 1500 as a rough guide, perhaps half or double the true figure.

COUNTING ENERGY SERVANTS

An energy servant is one HLE of physical work done by something other than a human body. To count them, we take the useful energy per person in a year and divide it by 0.684 GJ. Useful energy is what actually does work, after losses in burning and moving fuel, not the fuel itself.

The idea is not new. In 1940 the inventor Buckminster Fuller counted about 17 "energy slaves" for every person on Earth ^[23]. Our method is stricter, but the picture is similar. A Western European had about 2 energy servants in Rome's time, 5 in 1800, 25 in 1900 and 73 in 2000.

FIVE CHANNELS OF OUTPUT

Realised output is what an average person actually gets done in a year, with all their help. We build it from five channels and add them up. In 2026 the total is about 38 HLE per person, and machines supply about 37 of it.

1. **Own work.** Hours worked ^[25] and the share of people of working age ^[86]. Because children, old people and part-timers count, this is about half an HLE per person today.
2. **Machines.** Useful energy per person from motors, vehicles and other powered machines ^[27], turned into energy servants.
3. **Robots.** Robots per person ^[52], the hours they run, and how much work each does compared with a person.
4. **Digital agents.** Computer power per person in watts, and the energy for each accepted hour ^[50]. We add the price of AI work ^[40] and its time horizon ^[38]. That is the length of task, in human working hours, that an agent can finish on its own half the time.
5. **Virtual worlds.** The share of AI work and spending that goes on VR, AR and spatial computing. Think of immersive worlds and 3D digital twins. It is about 1% today. Music,

video, images and games on ordinary screens count as digital, not virtual.

THE COMMAND EQUATION

Realised output is what the average person gets. The Command Equation asks a different question: how much can one person direct when they choose to? It works like a recipe. You start with your own work, add what each helper does, then check which limit stops you first.

THE COMMAND EQUATION · 311 INSTITUTE

$$O = h \times [1 + \text{sum of } (n \times t \times s \times r) / T]$$

capped by the smallest of four limits: $h \times k$, B / p , E / e and D

O	output a person directs in a year, in HLE
h	the person's own working time (1 for a full-time adult)
sum of	add up over every channel: digital (software, content, services), physical (robots and machines), virtual (VR, AR and spatial worlds)
n, t	units directed in a channel and the hours each runs
s, r	speed relative to a human, and the share of work accepted after review
T	2,500 hours, one HLE
k	direction capacity: HLE of delegated work one person can specify and verify a year
B / p	production budget divided by the price of an HLE-year of delegated work
E / e	energy available divided by the energy per HLE-year
D	demand: what buyers, or the person, will actually absorb

Here is what each letter means in plain words. The letter O is the output a person directs in a year. The letter h is their own working time, which is 1 for a full-time adult. The letter n is the number of helpers in a channel, such as ten agents. The letter t is the hours each one runs in a year. The letter s is speed, or how many times faster than a person a helper works. The letter r is the share of its work that passes review. Dividing by 2,500 turns those helper hours into HLE.

The total can never be more than the smallest of four limits. In the box the limits are listed after the equation. Read them as a list of four: h times k, B over p, E over e, and D. Direction, h times k, is how much delegated work one person can specify, check and stand behind. Budget, B over p, is the money spent divided by the price of an HLE of work. Energy, E over e, is the power available divided by the energy for each HLE. Demand, D, is what buyers will actually take and pay for.

A WORKED CASE FOR 2026

Take a heavy user in October 2026: a skilled software builder on a flat \$180-a-month plan for a top AI agent service. They keep ten agents busy at the same time, for six hours a day, 300 days a year, which is 1,800 hours each. Each agent works five times faster than a skilled person, and half its work passes review. So ten times 1,800 times 5 times 0.5, divided by 2,500, gives about 18 HLE of accepted AI work.

TERM	VALUE	NOTES
n (agents)	10	concurrent frontier agents
t (hours each)	1,800	6 hours a day, 300 days
s (speed vs human)	5	cautious, tests suggest far higher before review
r (accepted)	0.5	half of agent output survives review
Delegated output	18 HLE	ten agents times 1,800 hours times speed 5 times acceptance 0.5, divided by 2,500 hours
Team equivalent	~18 people (up to ~45)	~\$2.7m a year at ~\$150k each, with all employment costs
Cost: subscription	\$2,160 a year	~\$180 a month
Cost: electricity	~1,800 kWh (~\$270)	~100 Wh per agent-hour
What limits it	Attention (k ~ 20)	budget, energy and demand do not bind for this heavy user

Source: 311 Institute analysis 2026 (worked case, our estimate).

That is the work of about 18 full-time people. It could reach about 45 if the agents are faster, or if more of their work is kept. Hiring 18 skilled people would cost about \$2.7m a year at \$150,000 a head. That is the fully loaded cost: salary plus the other costs of employing someone, such as tax, pensions and a desk.

A few words in the table need explaining. Frontier agents are the most capable AI agents on sale. Concurrent means running at the same time. GDPval is a test from OpenAI that scores AI on real job tasks ^[39]. It suggests agents can work far faster than five times a person, before review. The source line says "our estimate", which means the figures are our own working, not measured facts.

The agents use about 1,800 kWh of electricity in a year, worth about \$270. That is roughly what three or four fridge-freezers use in a year. For this heavy user, budget does not bind, since the plan is a flat fee. Energy does not bind either, and demand does not bind yet for a software builder with buyers.

What binds is direction. The builder has only so many hours to set tasks, check results and decide what comes next. Their direction capacity, k , is about 20 HLE a year, and at 18 they are already close to it.

SAME YEAR, TWO DIFFERENT LIMITS

This heavy user is not typical. The headline figure in the opening chapter, about 56 HLE in 2026, is for an ordinary skilled person. That person spends a tenth of an average income, about \$2,600 a year. Four-fifths of it buys about 17.6 HLE of AI work, at about \$120 for each HLE. Almost all of that is digital work. Only the virtual share, about 1% of the budget today, goes on VR, AR and spatial work. The last fifth buys robot work at about \$15,000 for each HLE, so it adds almost nothing yet. That is why the AI figure is 17.6, close to the heavy user's 18.

Their total is higher only because it counts more things. It adds their own work, which is 1 HLE, and their share of everyday machines, about 37.2 HLE. So 1 plus 17.6 plus 37.2 gives about 56. For this ordinary person the budget binds first, not attention. In the agentic diffusion branch the budget stays the limit for the ordinary person in every decade to 2100. Attention becomes the limit only for heavy users today, and in the explosive window from the 2030s.

FOUR BRANCHES AND THEIR ASSUMPTIONS

Our four branches share the same equation but use different inputs. Setback assumes a hard shock in the 2040s, such as a major war, a pandemic or a grid failure. It costs about a decade, then growth returns at 1.5% a year. Long drift assumes AI and robots spread slowly, with growth of 1.5% a year throughout.

Agentic diffusion assumes agents spread widely and get very cheap. Income per head grows 2.5–4.5% a year, then settles at 3%. By 2100 there is about one helper robot per person. Explosive window assumes AI can do most tasks for a time. Growth runs at 6–12% a year, then settles at 4%, with about five robots per person by 2100. The table sets out the main inputs side by side.

ASSUMPTION	SETBACK	LONG DRIFT	AGENTIC DIFFUSION	EXPLOSIVE WINDOW
Growth in income per head	Lost 2040s decade, then 1.5% a year	1.5% a year	2.5% to 4.5%, then 3%	6% to 12%, then 4%
Energy per head, 2100	~85 GJ	~108 GJ	~264 GJ	~554 GJ
Helper robots per person, 2100	0.05	0.1	1	5
Computer power per person, 2100	60 W	100 W	1 kW	3 kW

ASSUMPTION	SETBACK	LONG DRIFT	AGENTIC DIFFUSION	EXPLOSIVE WINDOW
Energy per accepted hour of AI work, 2100	8 Wh	5 Wh	0.5 Wh	0.05 Wh
Price of one HLE of digital work, 2100	\$20	\$10	\$1	\$0.10
Direction capacity k, 2100	~1,000	~3,000	~110,000	~2,000,000
Income per head, 2100 (2025 \$)	~\$53k	~\$79k	~\$373k	~\$3.0m

Source: 311 Institute analysis 2026 (our estimate), anchors from UN WPP 2024, IEA 2025, Davidson 2021, Epoch AI 2025, bank humanoid forecasts 2024–26.

WHERE THE ODDS STOP

We give a percentage chance only for events up to 2050, and every prediction in this report is settled by 2050. For a call more than ten years away we give a range, such as 25–45%, not one exact number. Each one rests on evidence we can point to today. Up to 2050, long drift marks the low edge of the band, agentic diffusion the middle, and explosive window the high edge.

After 2050 we give no odds at all, because nobody can honestly put a probability on 2100. Small changes in early years grow into huge gaps by the end of the century. Instead we show the four branches, the branch points that decide between them, and the decisions that pay off in every branch.

HOW WE SIZED EACH PREDICTION

Each probability in this report comes from our analysis, built in five steps. First, we start from base rates: how often similar changes happened in the past, and how fast. Second, we look at current trend data, such as prices, sales and test results. Third, we check physical limits, such as energy, materials and factory capacity.

Fourth, we apply a discount where a sector has a record of promising more than it delivers. Robot makers are the clearest case of this. A raw reading of their plans gives humanoid shipments of a million a year by 2032 about a 50% chance. We cut that to 30%. Fifth, we cross-check each call against the others, so the whole set fits together. If two calls depend on the same thing, their odds must agree.

The result is the figure printed as CONFIDENCE on each prediction block. It is the chance we give the call: how likely we think it is that the claim comes true. Each block also names an early sign, and a tripwire that tests the same thing as the claim.

HOW WE CHECKED THE SOURCES

In October 2026 we checked every figure in the report again against its origin. Where the origin had changed or we had read it wrongly, we corrected it. For example, METR, an independent lab that tests AI systems, measures each agent's time horizon ^[38]. METR first reported Claude Opus 4.6 at 14.5 hours, then rescored it at about 12 hours. We now use the lower figure, and the chart in the chapter on AI agents shows 12 hours. For tasks an agent must finish four times out of five, the horizon is about ten times shorter.

Each source is graded by trust tier: primary, corroborated, single-source, unverified or foresight. Where a claim is disputed, we flag it and give the range. Revenue that a company reports about itself, without an audit, is a common case. Two figures could not be checked again: energy per head in England around 1800 ^[10], and long-run output per hour from 1900 ^[24]. We keep them, but label them and soften the wording.

WHAT THE REPORT LEAVES OUT

This report measures output, not wellbeing, so it does not measure welfare, happiness or how much people consume. A person who directs a thousand times more work is not a thousand times happier. The report also treats the demand side only lightly. It asks what one person can make, not who will buy it all, though demand appears as one of the four limits.

MAIN LIMITS OF THE METHOD

The most important number in the report is also the least measured. That is direction capacity, k : how much delegated work one person can specify, check and stand behind. Today we can only estimate it from what heavy users report. If k grows slowly, every one of the high branches shrinks with it.

Forecasts from banks and robot makers carry a second risk. These firms have commercial interests in the outcome, so we use their numbers with care and say where they come from. Third, many AI company revenues are run-rates (this month's sales times twelve). They are not audited, so they can overstate a firm's real size.

Last, the HLE adds up things that are not alike. An accepted hour of code and an hour of a robot stacking boxes both count as one hour of work. That is a choice, not a law of nature. We make it because it shows what a person could otherwise have hired. Readers who would make different choices can use the data book to do so.

I CHAPTER SUMMARY

We count work in HLE and money in 2025 dollars. History runs in 100-year steps through Rome and Western Europe, then 10-year steps for an average global family. Output per person is built from five channels and capped by four limits. Odds stop at 2050, and four branches follow. Every source was checked again in October 2026.

Check the method first, and only then judge the numbers.

CHAPTER CLOSE

WHAT NEXT?

How we built the numbers

QUESTIONS TO ASK

- 1 Which assumption in this chapter would you change first?
- 2 How would you measure your own direction capacity?

03

PART THREE

ONE BODY, ONE OUTPUT

Before Rome: the lone human body and the first ways to multiply it.

03

How much could one human body make before anyone owned a tool, an animal or another person?

For almost all of human history, one person made about one person's worth of output. The story of progress is the story of breaking that ratio.

SCOPE

300,000 BC

Farming

First states

Every story about the future of work starts with a quiet idea, so quiet that we rarely notice it. One person sells one person's work. A carpenter sells hours of carpentry, and a coder sells hours of code. The price changes with skill, scarcity and place, but the unit does not. One body, one output. This report asks what happens when that unit breaks. To answer, we go back to the beginning, because the ratio has broken before, and each time it broke, it changed who got rich. Now AI agents and robots are starting to break it again. To see how big that break could be, we first need to know how much one person could make on their own.

Three words matter all the way through. **Output** is how much gets made, and **productivity** is output for each hour of work. **Capture** is who keeps the value. They sound alike, but they are not the same. A slave plantation could have huge output and high productivity, while the people who did the work captured almost nothing. This report keeps the three apart.

ONE BODY ON ITS OWN

Engineers who measure human power agree on a clear number. A healthy adult can keep up about 75 watts of useful physical work across an eight-hour day ^[8]. A watt is a unit of power, the rate at which work gets done. 75 watts is about a tenth of a horsepower, or the power of one old-style light bulb. In short bursts a person can reach 300–400 watts. Over a working year of about 2,500 hours, 75 watts adds up to around 190 kilowatt-hours (kWh) of work. That is about what a 2-kilowatt electric heater uses if it runs for four days without stopping. It is the physical budget of a whole year of human labour. This report calls it one **Human Labour Equivalent**, or HLE.

An HLE is counted in one of two ways only. The first way is physical: about 190 kWh, or 0.684 gigajoules (GJ), of useful work. A gigajoule is a billion joules, a unit of energy. The second way is time: 2,500 hours of accepted work. That work can be done by a person, a robot or an AI agent. An AI agent is software that can carry out a task on its own, step by step. 'Accepted' means someone checked the work and kept it.

The HLE never measures money. Value is always counted on its own, in US dollars at 2025 prices. For example, a Roman-era adult farming with hand tools might grow about 800 kilograms of wheat-equivalent in a year. Wheat-equivalent means wheat, or other food worth the same. At today's prices that food is worth about \$1,000. The chapter on method sets out both ways of counting the HLE. The point here is simpler. For roughly 300,000 years, every human started with one HLE and nothing else.

PRIME MOVER	SUSTAINED POWER	HLE (POWER)	TRUST TIER
Adult human, 8-hour day	~75 W (50–100)	1	Corroborated
Donkey	~150–250 W	2–3	Our estimate
Ox	~300–400 W	4–5	Corroborated
Horse	~500–750 W	7–10	Corroborated
Roman or medieval watermill	~1–4 kW	15–50	Corroborated
Dutch windmill, 17th century	~5–10 hp	50–100	Our estimate
Newcomen engine, 1712	~5–15 hp	50–150	Corroborated

PRIME MOVER	SUSTAINED POWER	HLE (POWER)	TRUST TIER
Boulton & Watt engine, c.1800	~15–20 hp	150–200 (450–600 round the clock)	Corroborated
Large mill engine or locomotive, 1900	100–1,000+ hp	1,000–10,000+	Our estimate

Source: Smil 2017, Kanefsky & Robey 1980, 311 Institute analysis 2026 (historical baseline). HLE on a power basis: 75 W kept up over a 2,500-hour working year.

NEANDERTHALS AND THE ONE-TO-ONE WORLD

The first modern humans, *Homo sapiens*, appeared around 300,000 years ago. Neanderthals lived at the same time. Both lived in bands of 25–50 people, and both had stone tools, spears and fire. Fire was the first source of outside energy, and people have used it under control for at least 400,000 years. It cooked food, which freed calories the gut could not reach on its own, and it kept the band warm at night. The historian Ian Morris finds that foragers used about 4,000–5,000 kilocalories a day per person ^[9]. A kilocalorie is the unit printed on food labels. About 2,000 of those came from food, and the rest came from fuel, clothing and shelter.

What foragers could not do was *store* an advantage. Bands shared what they had, so there was little private surplus (anything beyond what people needed to live). A hunter who killed a mammoth did not own a mammoth's worth of output, because the band ate it. Output and use ran almost one to one. Tools did exist, of course. The hand axe, the spear-thrower and later the bow all raised what a body could do. But none of them added a second body's worth of work. That ratio of roughly one to one is the starting line for everything in this report.

FARMING: MORE OUTPUT, WORSE HOURS

From about 10,000 BC, people began to farm. It started in the Fertile Crescent, the arc of land from Iraq to Egypt. It also began on its own in China and in Mexico and Central America. Output per hectare (a field about 100 metres by 100 metres) rose perhaps ten to a hundred times. Population followed. But output per *hour* almost certainly fell. Farmers worked longer than foragers, and their skeletons show that they grew shorter too. The economist Gregory Clark goes further. He argues that average living standards in 1800 were no better than in the Stone Age ^[17]. In his view, every gain went into more people, not richer ones. Historians argue about how big that effect was, but most agree on its direction.

Farming matters to this story for another reason. It created **surplus that could be taken**. Grain can be stored, taxed and stolen, and land can be owned. Once surplus existed, a new kind of multiplier could arise. One person could now live from the output of many.

THE FIRST TRUE MULTIPLIERS

Between about 4,000 and 3,000 BC, three things arrived at once. First came draft animals pulling the ard, a simple scratch plough. An ox gives 300–400 watts, the power of four or five people, to whoever owns it. Second came the wheel and the first states, such as Uruk in Iraq and early Egypt, and states added roads, writing and taxes. Third, slavery became an institution that turned human bodies into property. For the first time, one person could command the output of others: animals, slaves and farmers who paid tax.

The pattern that runs to 1900. From the first states until the steam age, there was one main way to multiply a person's output beyond about one HLE. You had to own or command other bodies, such as oxen, horses, slaves, serfs (farm workers tied to a lord's land) and tenants. When machines came, they moved the multiplier from bodies to capital, meaning money put into machines, buildings and tools. Ownership stayed in few hands, and this report treats forced labour honestly, as *other people's multipliers*. The output gains were real, but the people who paid for them did not keep them.

The next five chapters follow that pattern through two thousand years. They start from a Roman family of four in 1 AD. A typical smallholder family then did about 2.5 HLE of work a year. The average family, counting the rich, did about 10 HLE. Our long series follows that average family, and then switches to the average family on Earth today.

CHAPTER SUMMARY

On its own, a human body gives about 75 watts of steady work, or about one HLE a year. Fire, farming, animals and states each raised what a person could direct. But before machines, there was only one way to multiply one person's output many times over. You had to command other people and their animals.

The first multiplier was other bodies, and the few owned it.

CHAPTER CLOSE

WHAT NEXT?

One body, one output

PREDICTIONS

PREDICTION

30% CONFIDENCE · BY 2035

By 2035 a G7 statistics office, the ILO or an employer with more than 10,000 staff publishes a regular measure of delegated output per worker. That means work done by AI agents or robots that a person directs.

BASIS

Counting hours misreads AI-assisted work. In a 2025 trial by METR, an independent AI testing lab, coders were 19% slower with AI but felt faster. The G7 are seven large rich democracies. The ILO is the International Labour Organization. Early sign: a statistics office asks how much work AI agents complete.

TRIPWIRE

Proved wrong if no G7 statistics office, the ILO or an employer with over 10,000 staff publishes one by 2035.

QUESTIONS TO ASK

- 1 Which of today's multipliers are owned by the few, as oxen once were?
- 2 Which new tools raise total output but make each hour worse?
- 3 What does one person working alone produce in your work?

04

PART FOUR

THE ROMAN FAMILY OF FOUR

The baseline: what a Roman household of four made in 1 AD.

04

What did a Roman family make in a year, and who held the multipliers?

A typical Roman smallholder family of four did about 2.5 HLE of work a year, worth roughly \$2,500. The average family, counting the rich, did about 10 HLE, worth about \$9,900. One of Rome's richest freed slaves owned 4,116 slaves and 3,600 pairs of oxen. He directed about 33,000 HLE, about 13,000 times the smallholder family.

SCOPE

Rome 1 AD

Slavery

Average vs typical

To measure two thousand years of change, we need a unit that stays put. This report uses a Roman family of four in 1 AD, with two adults and two children in Italy. But there are two Roman families to keep apart. The **typical** family worked a small farm and did about 2.5 HLE of work a year. The **average** family is what you get if you share all of Rome's output equally. It did about 10 HLE, because the rich pulled the average up. Our long series follows the average family, century by century, to 2000. After 2000 we switch to an average global family that is not tied to any one country. By then, Europe is no longer the place where output per person is highest.

Unit

Family of four

Year

1 AD

Typical family

~2.5 HLE

Average family

~10 HLE

THE EMPIRE IN NUMBERS

The emperor Augustus ruled perhaps 45–60 million people, and by 150 AD the empire held about 70 million. To compare wealth, economists use GDP per head. GDP is the value of everything a country makes in a year, and 'per head' divides it by the number of people. Historians often state it in dollars at 1990 prices, a fixed yardstick that removes the effect of inflation. Estimates for Rome differ widely, and the size of that range is itself a finding. Angus Maddison put Roman Italy at about \$809 a head on that basis. Elio Lo Cascio and Paolo Malanima put it far higher, at \$1,300–1,800, which is close to Italy's level between 1500 and 1800 ^[2].

Walter Scheidel and Steven Friesen measure output in another way, in food. They estimate the whole empire's output at 620–680 kilograms of wheat-equivalent per person each year ^[3]. That is about 1.7–2 times bare subsistence, the least a person needs to stay alive. Three-quarters or more of all workers farmed the land.

Rome also used a lot of energy for a world without machines. Morris's index puts the Roman West at about 31,000 kilocalories per person per day ^[9]. The West did not reach that level again until about 1700. Much of it was hay and grain for animals, and wood for heat. Little of it reached the small farmer as useful work.

THE HOUSEHOLD LEDGER

A typical small farming family worked two to five hectares. It usually had no ox of its own and perhaps shared one with neighbours. It needed about 800–1,000 kilograms of wheat-equivalent just to eat. Once clothing, fuel, seed and dues were counted, it needed 1,300–1,600. In a good year it grew about 1,500–2,500 kilograms. The two adults worked 200–250 days a year, for eight to ten hours a day. That came to 4,000–5,000 hours between them. Children helped from about age seven. Grinding grain on a hand quern, a pair of round stones turned by hand, took two to three hours of a woman's day.

ROMAN SMALLHOLDER FAMILY, 1 AD		VALUE	TRUST TIER
Household	2 adults, 2 children, 2–5 ha		Our estimate
Food need	~800–1,000 kg wheat-eq a year		Our estimate
Total need (with clothing, fuel, seed, dues)	~1,300–1,600 kg		Our estimate
Gross output	~1,500–2,500 kg wheat-equivalent (about \$2,500)		Our estimate
Work done	~2.5 HLE (two adults, children, a shared ox)		Our estimate

ROMAN SMALLHOLDER FAMILY, 1 AD		VALUE	TRUST TIER
Paid in rent, tax or interest		~10–30%	Our estimate
Hours worked		~3,200–5,000 adult hours	Our estimate
Multipliers	Hand quern, a shared ox (~4 HLE part-time)		Our estimate

Source: Hopkins 1980, Scheidel & Friesen 2009, Columella (1st century AD), 311 Institute analysis 2026 (historical baseline).

In this report's units, the typical smallholder family did about **2.5 HLE** of work a year. That came from two adults, some help from the children and a part-time share of an ox. In money, that work was worth roughly **\$2,500** a year at today's prices. A free, skilled craftsman in a town did better than this. Diocletian's Price Edict of 301 AD set a carpenter's wage at 50 denarii a day plus food ^[7]. That was double a farm worker's 25 denarii. Skill doubled the price of an hour, but it did not change the unit.

AVERAGE FAMILY VERSUS TYPICAL FAMILY

Here is the first surprise in the data. Divide Roman Italy's output by its number of people. The *average* family of four then comes out at about **\$9,900** a year, with about 10 HLE of useful work. That is four times what our farming family made. The gap is not a mistake in the numbers, but comes from how output was shared out. Scheidel and Friesen estimate that the top 1.5% took about a fifth of all income. These were the senators, rich knights and town councillors. The same study finds that nine in ten Romans lived at or near subsistence ^[3]. A few households that directed huge multipliers drag the average up. This average family is the one our long series follows from here to 2000.

THE MULTIPLIER OWNED BY THE FEW

Consider Gaius Caecilius Isidorus, a freed slave who died in 8 BC. The writer Pliny records that his will listed 4,116 slaves, 3,600 pairs of oxen and 257,000 other animals ^[5]. Count each slave as one HLE and each pair of oxen as about eight. Then Isidorus directed some **33,000 HLE** of work a year. Most of it came from the oxen, about 29,000 HLE, and about 4,100 from the slaves. That is about 13,000 times the work of a typical smallholder family. In Rome itself, the city prefect Pedanius Secundus kept about 400 slaves in his town household. All of them were executed after his murder in 61 AD ^[6].

Slavery was the empire's largest single multiplier. Scheidel puts the number at one to one-and-a-half million slaves in Italy ^[4]. That was 15–25% of all people there, and about a tenth across the whole empire. Watermills existed too. The mill complex at Barbegal in Gaul, now southern France, had 16 waterwheels. Older estimates said it could grind flour for about ten thousand people. A 2018 study

suggests it made mostly ship's biscuit and stood idle for months each year^[104]. Mills were rare, and estates and cities owned them, so the family of four saw almost none of their output.

WHO	WHAT THEY DIRECTED	HLE A YEAR	RATIO TO FAMILY
Smallholder family of four	Two adults, children, share of an ox	~2.5	1
Average Roman Italian family (mean)	Includes elite income	~10	~4
Pedanius Secundus, urban household	~400 slaves	~400	~160
Caecilius Isidorus (d. 8 BC)	4,116 slaves, 3,600 pairs of oxen	~33,000	~13,000

Source: Pliny the Elder, Natural History 33.135, Tacitus, Annals 14, 311 Institute estimate 2026 (historical baseline). Slaves counted at ~1 HLE, a pair of oxen at ~8 HLE.

Hold on to that ratio of about 13,000 to 1, the gap between one of the most powerful private commanders of output and an ordinary family. Later chapters look at 2100 through four branches, which are paths that could happen, with no odds attached to them. In the agentic diffusion branch, a skilled person spending a tenth of an average income directs about 30,000 HLE a year. That is roughly what Isidorus directed. In the other branches it is about 260 (setback), 700 (long drift) or 2 million (explosive window). On every branch, the big difference from Rome is that nobody will be owned.

CHAPTER SUMMARY

A typical smallholder family grew about 1,500–2,500 kg of wheat-equivalent a year. It did about 2.5 HLE of work, worth roughly \$2,500, and handed 10–30% to rent, tax or lenders. The average family, counting the rich, did about 10 HLE, worth about \$9,900, because the top 1.5% took a fifth of all income. Our long series follows the average family.

Rome's multipliers were real, but almost none belonged to the family that did the work.

CHAPTER CLOSE

WHAT NEXT?

The Roman family of four

PREDICTIONS

PREDICTION

5–15% CONFIDENCE · BY 2040

By 2040 one person with no employees directs AI and robot work equal to 10,000 full-time staff. That would be Isidorus-scale command, built from rented agents and robots.

BASIS

Our analysis puts a frontier operator, meaning a skilled person backed by plenty of money, at only about 300–700 HLE a year by 2040. So this call needs someone 15–30 times above that path. Early sign: by 2032 a solo operator shows a credible audit of more than 1,000 HLE of accepted AI work in a year.

TRIPWIRE

Proved wrong if no credible audit shows one person with no employees directing 10,000 HLE or more by 2040.

QUESTIONS TO ASK

- 1 What is today's 'shared ox', a multiplier households borrow but do not own?
- 2 Which modern fortunes rest on controlling the multiplier, not doing the work?
- 3 What if your organisation's 'Isidorus' assets were rented to everyone?

05

PART FIVE

MILLS, MONKS AND THE LONG FLAT LINE

From 1 AD to 1500: the frontier moves, but the ceiling does not.

05

Why did fifteen centuries of invention leave the ordinary family making about the same?

Output per head stayed within a band of about two to three times for 1,500 years. Watermills, windmills, ploughs and printing raised the ceiling, but mostly for the people who owned them.

SCOPE

Fall of Rome

Medieval Europe

Song China


If the Roman family of four could see its descendants a thousand years later, it would not be impressed. Around 1000 AD, Western Europe's output per head stood at about \$427 in Maddison's 1990 prices ^[1]. That was well below Roman Italy, and the frontier, the place where output per person was highest, had moved east. The Song dynasty ruled China from 960 to 1279. On one careful estimate, Song China reached about \$1,250–1,500 a head in the same 1990 prices around 980–1020 AD ^[102]. By one estimate, its iron industry produced more than 100,000 tonnes a year by 1078 ^[103]. In Western Europe, the family of four made about **\$3,800** a year in 2025 money, and it did about 8 HLE of useful work.

Family of four, 1 AD to 1500: the long flat line

SOURCE: MADDISON, LO CASCIO & MALANIMA, 311 INSTITUTE ANALYSIS 2026 (HISTORICAL BASELINE)


VALUE, \$ THOUSAND (2025 PRICES)

\$ thousand a year



USEFUL WORK, HLE

HLE a year



COLLAPSE AND RECOVERY

Three blows hit after Rome. The Western empire fell, the Justinianic plague struck from 541, and then came centuries of small, warring kingdoms. Together they cut output per head by more than

half. Trade shrank, and cities emptied. Country estates turned inward and fed themselves. By the 600s and 700s the European family made only about \$3,900 of value a year. The average Roman family at its peak had made about \$9,900. The multipliers did not vanish, because oxen still pulled ploughs, but what vanished were the networks that let surplus travel.

THE MILL AS MONOPOLY

The great machine of medieval Europe was the watermill. In 1086 William the Conqueror ordered a survey of England, the Domesday Book. It counted about 5,624 mills, roughly one for every 200–300 people ^[11]. Each mill produced perhaps two horsepower, about 1.5 kilowatts, the work of about 20 people. In all, England's mills supplied about 110,000 HLE of grinding power. With under two million people, that was about 0.06 HLE per person.

Lords and abbeys owned those mills. They held rights called *banalités*, French for a lord's right to run the only local mill, oven or wine press. Under these rights they forced tenants to use the lord's mill and pay a fee. Hand querns were sometimes taken away so that peasants could not grind grain at home. The pattern is exact. A machine multiplies output. A few people own it, many people rent it, and the owner sets the rent. This pattern will come back in every chapter of this report.

PLOUGHS, COLLARS AND WINDMILLS

Between the ninth and twelfth centuries, Europe took up three farm advances. The heavy plough could turn thick, wet soil. Three-field rotation left only a third of the land resting each year, not half. The padded horse collar let a horse pull without choking. A horse has the power of seven to ten people, and the collar made it a practical farm animal. Windmills appeared in the 1180s. These were all real gains. By 1300, England's output per head had climbed to about \$755 ^[1]. On one estimate, Italy's trading cities, such as Venice and Florence, reached about \$1,400–1,650 ^[105].

THE BLACK DEATH DIVIDEND

Then, between 1347 and 1351, plague killed between a third and a half of Europe's people. This was the Black Death. The survivors inherited land, animals and tools, and because workers became scarce, wages rose. Serfdom collapsed in Western Europe. Output per head rose, but not because of any new machine. The same stock of multipliers was simply shared among fewer hands. This is one of history's few clear cases of the family capturing more of the output, because a shortage of workers gave them bargaining power.

PRINTING: THE FIRST INFORMATION MULTIPLIER

Around 1450, Gutenberg's printing press changed another kind of output. A scribe could copy a few pages a day, while a press could print about 3,600. In the last half of the fifteenth century, Europe printed perhaps 12.6 million books ^[12]. Over that whole century, scribes made only about five million handwritten books. This was the first great multiplier of *information* output. It is the closest match in history to the AI agents in later chapters. The profits went to printers and merchants, and to rulers who licensed them. Readers got books at a fraction of the old price, a gain that shows up as falling prices, not as higher measured output.

The table below follows Western Europe's output and family value. Its 'Top multiplier' and 'Who held it' columns name the leading machine or system of each age and who controlled it. In 500 that was the Byzantine state, in the East, and the Church.

YEAR	GDP A HEAD (1990 \$)	FAMILY VALUE (2025 \$)	TOP MULTIPLIER	WHO HELD IT
1	1,100	\$9,900	Slaves, oxen, roads, shipping, first watermills	Senatorial and equestrian landowners. The state
500	550	\$4,950	Monastic watermills	Kings, monasteries and the Church
800	450	\$4,050	Heavy plough, three-field rotation	Manorial lords. The Church
1000	427	\$3,843	Watermills: 1 per ~200–300 people by 1086	Lords and abbeyes (milling monopolies)
1300	750	\$6,750	Water-powered fulling and sawing. Compass	Merchant republics. Lords
1400	790	\$7,110	Labour scarcity after the Black Death	Surviving peasants gain
1500	771	\$6,939	Printing press. Caravel	Printers, merchants, Crown monopolies

Source: Maddison 2007, Lo Cascio & Malanima 2009, Broadberry et al. 2015, 311 Institute analysis 2026 (historical baseline).

By 1500 the European family made about **\$6,900** a year in value. It did about 11 HLE of useful work. Its useful work was back around Rome's level, though its value was still about a third lower than the Roman average. Fifteen centuries of invention had moved the frontier from Rome to Song China, then to Italy and on towards Holland. But they had not broken the ceiling, and that needed a new source of energy.

CHAPTER SUMMARY

After Rome, the European family's output per head fell by more than half. Useful work took about a thousand years to recover, and value had still not caught up by 1500. Mills and better ploughs added useful work. But lords and abbeys owned the mills and charged for their use. For long stretches, Song China led the way.

New machines raised output, but their owners kept most of the gain.

CHAPTER CLOSE

WHAT NEXT?

Mills, monks and the long flat line

PREDICTIONS

PREDICTION

70% CONFIDENCE · BY 2032

By 2032 a major competition authority opens a formal market investigation into access to frontier AI models or compute. The authority could be in the US, the EU, the UK, China, Japan or India.

BASIS

A competition authority polices monopolies. Frontier models are the most capable AI. Compute is computer processing power. A few firms control both. A binding price rule is less likely, about 15–20%. Early sign: a market study of AI chips or cloud.

TRIPWIRE

Proved wrong if no such authority opens a formal market investigation into AI model or compute access by the end of 2032.

QUESTIONS TO ASK

- 1 Which of today's 'mills' charge a toll for access, and who sets it?
- 2 Where is today's frontier economy, and are you watching it?
- 3 Would a sudden labour shortage shift who captures AI's output?

06

PART SIX

SHIPS, COAL AND STEAM

1500 to 1900: the ceiling breaks, first for owners, then for workers.

06

When did one family's output finally escape the band, and what did it cost?

Atlantic slavery, coal and steam raised output per head nearly fourfold in Western Europe. Workers' real wages rose only after about 1830.

SCOPE

Atlantic world

Industrial Britain

Asia diverges

After 1500 the frontier moved to the Atlantic. By 1600, Holland's output per head reached about \$2,600 in 1990 prices, the highest in the world ^[1]. It rested on windmills, cheap ships and the first joint-stock companies. The Dutch alone ran perhaps 8,000–10,000 windmills of five to ten horsepower each. From 1595 the fluyt, a Dutch cargo ship, needed fewer crew per tonne than any ship before it. The Dutch East India Company (VOC), founded in 1602, was a joint-stock company, so many shareholders could own part of it without working in it. In this report's terms, they owned a share of the multiplier without working it.

OTHER PEOPLE'S MULTIPLIERS

The same century built the largest system of forced labour in history. Between 1501 and 1866, about 12.5 million Africans were forced onto ships bound for the Americas ^[14]. About 10.7 million of them survived the crossing. By 1860 the United States held about 3.95 million enslaved people ^[15]. Their market value was commonly estimated to be more than that of all the country's railways and factories combined. In Eastern Europe, a 'second serfdom' tied peasants to the land, and Russia freed about 23 million privately owned serfs only in 1861 ^[16].

In output terms, a planter or a landowner held multipliers of a hundred to ten thousand HLE. They were built entirely from other people's bodies. When Britain abolished slavery in 1833–34, it paid £20 million in compensation, roughly US\$3 billion at today's prices by our estimate. The money went to the owners, not to the people who had been enslaved. This history is not a footnote to the story of output. It is the clearest proof of this report's central point: **output, productivity and capture are different things**. A plantation could be very productive and its output very large, yet the people who did the work captured almost nothing.

COAL AND STEAM POWER


Coal changed the sums. England burned about three million tonnes around 1700. In 1712, Thomas Newcomen's steam pump began lifting water out of mines, with five to fifteen horsepower. James Watt patented a better steam design in 1769 and sold it from 1776. His machines averaged fifteen to twenty horsepower. That is the work of 150–200 people, or three times that if the machine ran day and night. By 1800 Britain had about 35,000 horsepower of steam and 120,000 of water power ^[13]. That was perhaps 1.5–2 million HLE of fixed power, for a workforce of under five million.

By 1870, Britain's steam power had reached about 2.06 million horsepower ^[13]. That was around 20 million HLE on a day-shift basis, or 60 million if it all ran nonstop. The workforce was about twelve million, and railways and steamships came on top. By one estimate, energy use per person in England and Wales reached about 55–65 gigajoules by 1800 ^[10]. That was about three times Italy's level, and on the same rough estimate it roughly doubled again by 1900.

Muscles and machines compared

SOURCE: SMIL 2017, KANEFSKY & ROBEY 1980, 311 INSTITUTE ANALYSIS 2026 (HISTORICAL BASELINE)

POWER OF ONE UNIT, HLE, LOG SCALE



WHO GOT THE STEAM

For the first two generations of factories, mill owners and other capitalists captured most of the gain. Working hours rose. The historian Jan de Vries calls this an 'industrious revolution', as English working years passed 3,000 hours. Real wages, meaning pay adjusted for prices, barely rose until about 1830–1850. The economist Robert Allen measured the gap. Between 1780 and 1840, British output per worker rose 46%, but real wages rose only 12% ^[96]. He called this Engels' pause, after Friedrich Engels, who wrote about the factory poor.

Only after that pause did workers' pay begin to rise steadily. Unions, factory laws and the sheer scale of new output all helped. Over time the labour share, the part of a nation's income paid to workers, rose too. The lag matters for the AI chapters later in this report. It suggests that even a technology that lifts everyone in the end can concentrate its first decades of gains.

THE GREAT DIVERGENCE

By 1900 the United Kingdom produced about \$4,490 a head in 1990 prices. The United States made about \$4,090 and Western Europe about \$2,890 ^[1]. China stood near \$550 and India near \$600, so a gap of about eight to one had opened between the leaders and the rest. Historians call this split the great divergence. The European family of four now made about **\$26,000** a year in value. It directed about 102 HLE of useful work, mostly from coal. Yet its own labour was still only about 2.5 HLE across its four members.

YEAR	GDP A HEAD (1990 \$)	FAMILY VALUE (2025 \$)	FAMILY WORK (HLE)	TOP MULTIPLIER
1500	771	\$6,939	11	Printing press. Caravel
1600	894	\$8,046	13	Dutch windmills. Fluyt. Joint-stock companies
1700	997	\$8,973	15	Coal. Newcomen engine (1712)
1800	1,202	\$10,818	23	Watt engine. Factories. Coke iron
1900	2,890	\$26,010	102	Steam, rail, steel, electricity, chemicals

Source: Maddison, Broadberry, Kander, Malanima & Warde 2013, Kanefsky, 311 Institute analysis 2026 (historical baseline).

CHAPTER SUMMARY

Between 1500 and 1900, the European family's yearly value rose about fourfold, from about \$6,900 to \$26,000. Its useful work rose about ninefold. The first gains rode on forced labour and on coal. Steam then moved the multiplier from bodies to machines. But machine owners captured most of the gain for two generations.

Machines replaced bodies as the multiplier, but ownership did not change hands.

CHAPTER CLOSE

WHAT NEXT?

Ships, coal and steam

PREDICTIONS

PREDICTION

45% CONFIDENCE · BY 2032

From 2026 to 2032, US output per worker grows at least 0.6 percentage points a year faster than the median worker's real wage. An example would be 2.0% against 1.4%. The median worker is the one in the exact middle of the pay scale.

BASIS

History shows that new gains can reach owners first. Robert Allen found that in Britain from 1780 to 1840, output per worker rose 46% while real wages rose only 12%. Labour's share of income has drifted down since 1980, AI tools are spreading fast, and a few firms own most compute. Early sign: US data for 2027–28 show output per worker growing well ahead of median pay, and the labour share falling.

TRIPWIRE

Proved wrong if output per worker outgrows median real wages by less than 0.6 points a year.

QUESTIONS TO ASK

- 1 How long did steam's gains take to reach wages? Is AI faster?
- 2 Which of today's supply chains still rest on other people's multipliers?
- 3 What is the coal of the AI age, whose owners gain first?

07

PART SEVEN

THE CENTURY OF ENERGY SERVANTS

1900 to 2000: the family directs hundreds of machine workers, and wages keep up.

07

How did the twentieth century hand an ordinary family hundreds of invisible servants?

Electricity, motors, the assembly line, the shipping container and the computer took the European family from about 100 to about 290 HLE of useful work. For once, wages kept up.

SCOPE

Electricity

Mass production

Computers

In 1940 Buckminster Fuller, an American inventor, wrote in *Fortune* magazine about 'energy slaves' ^[23]. He meant the fossil energy doing work for each person. He counted about 17 for every person on Earth, and Americans held more than half of the world's total. By 1950 he had raised the world average to 38. The phrase is harsh, and Fuller meant it to be. Fuller's point was that coal, oil and gas, piped into motors and machines, now did the physical work that slaves and serfs once did. And they did it for everyone, not only for owners.

FIVE MULTIPLIERS THAT CHANGED THE CENTURY

Electricity reached factories slowly. Electric motors supplied about 5% of US factory power in 1900, half by 1920 and four-fifths by 1929. The gain in productivity came only in the 1920s, about forty years after the dynamo, an early electric generator. It came once factories were rebuilt so that each machine had its own motor. **The moving assembly line** cut the time to build a Model T chassis from 12.5 hours to about 93 minutes in 1913–14 ^[18]. The chassis is a car's frame and running parts. **Tractors** helped cut US farm jobs from about 41% of all workers in 1900 to under 2% in 2000 ^[19]. They also replaced a working stock of more than 25 million horses and mules, counted in 1920.

The shipping container arrived in 1956 and cut the cost of loading cargo from about \$5.86 a ton to about 16 cents ^[20]. Research by Bernhofen and others found that dock work rose from about 1.7 to about 30 tons per hour of labour. And **the computer** raised computing performance by between 1.7 trillion and 76 trillion times, from hand calculation to 2006 ^[21]. Most of that rise came after 1945. Each of these five advances multiplied what one worker could produce. Each also, in the end, made goods so cheap that their measured value understated their volume.


HOURS FELL, OUTPUT SOARED

In 1870 a full-time worker in Western Europe or the United States worked about 3,100 hours a year ^[25]. By 2000 the Western European figure had fallen to about 1,350–1,800. Over the same period, output per person rose about six times and energy per person about five times. So output per hour rose by about ten to fifteen times. One long-run database, built by Bergeaud, Cette and Lecat, gives a similar picture ^[24]. By its estimate, US output per hour rose about eleven times from 1900 to 2025, Britain's about nine times and France's about eighteen times.


Family of four, 1500 to 2000: the break

SOURCE: MADDISON, KANDER, MALANIMA & WARDE, 311 INSTITUTE ANALYSIS 2026 (HISTORICAL BASELINE)

VALUE, \$ THOUSAND (2025 PRICES)



USEFUL WORK, HLE



WORKERS KEPT THEIR SHARE, UNTIL 1980

Unlike the steam age, the twentieth century's multipliers were largely captured by workers. Real wages rose with productivity, and hours fell. The labour share stayed broadly steady until about 1980. Unions, mass schooling and tight job markets after two world wars all played a part. So did taxes that took a bigger share from higher incomes. Since about 1980, the share has drifted down. Across many countries it fell by about five percentage points ^[31]. Worldwide it fell by 1.6 points between 2004 and 2024, to 52.3% ^[30]. A person's output can rise without their income rising in step.

THE LIGHT THAT BECAME FREE

William Nordhaus's history of lighting shows why output volume and output value split apart. Light is measured in lumen-hours, and 1,000 lumen-hours is about what one bright modern bulb gives off in an hour. A Babylonian lamp burning sesame oil needed about 41.5 hours of work to pay for 1,000 lumen-hours of light. A tallow candle in 1800 needed about 5.4 hours, and by the 1990s it took a fraction of a second ^[22]. GDP hardly noticed the change, because the price fell as fast as the volume rose. This gap between volume and value is the most important idea for reading the chapters on AI.

MULTIPLIER	PERIOD	EFFECT
Electric unit drive	1900–1930	5% to ~80% of US factory power, productivity payoff ~40 years after the dynamo
Moving assembly line	1913–14	Model T chassis 12.5 hours to ~93 minutes (~8x)
Tractors and engines	1900–2000	US farm workforce ~41% to ~1.9%, more than 25m horses and mules replaced

MULTIPLIER	PERIOD	EFFECT
Shipping container	1956 on	Loading \$5.86 to about \$0.16 a ton, dock work per labour-hour up many times
Computers	1945–2006	Computing performance up 1.7 to 76 trillion-fold since manual calculation
Electric light	1800–1990s	Labour hours per 1,000 lumen-hours 5.4 to ~0.0001

Source: David, Henry Ford Museum, USDA ERS 2005, Levinson 2006, Bernhofen et al. 2013, Nordhaus 1996, 2007 (historical baseline).

By 2000 the Western European family of four made about **\$175,000** a year in value. It directed about 294 HLE of useful work. Its own labour had shrunk to about two HLE, because its four people worked fewer hours, studied longer and retired earlier. More than 99% of the family's useful work was done by machines. The ratio had broken for muscle, but not yet for thinking.

The gap to watch. Our analysis expects the gap between volume and value to grow much wider in the AI era. After 2050 this report follows four branches, which are paths that could happen, with no odds attached. In the agentic diffusion branch, output volume per person rises about 100 times between 2026 and 2100. But measured income per head rises only about 14 times. So measured value understates output volume by about 1.7 times by 2060 and about 7 times by 2100. In the setback and long drift branches the gap is smaller, and in the explosive window branch it is larger.

CHAPTER SUMMARY

Between 1870 and 2000, output per hour in rich economies rose about ten to fifteen times. Working hours fell by nearly half. The Western European family came to direct about 290 HLE, almost all of it machine work powered by fossil fuels. Labour's share of income held steady until about 1980. It has fallen since.

For once the multiplier and the wage rose together, until about 1980.

CHAPTER CLOSE

WHAT NEXT?

The century of energy servants

PREDICTIONS

PREDICTION

40% CONFIDENCE • BY 2035

From 2026 to 2035, an official price index for AI or cloud computing services falls by at least 20% a year on average. A price index is a government measure of how prices change over time.

BASIS

Market prices for AI are falling much faster than this. Epoch AI, a research group, finds the price of a fixed level of AI ability falls 9–900 times a year, depending on the task. But official price indexes adjust slowly, often lag the market and mix AI with older services. Early sign: a national statistics office publishes a separate price index for AI services by 2028.

TRIPWIRE

Proved wrong if no official AI or cloud price index shows an average fall of 20% or more a year over 2026–2035.

QUESTIONS TO ASK

- 1 What made the 1900–1980 deal possible, and does it still hold?
- 2 Which of your firm's multipliers never show up in its accounts?
- 3 Which past multiplier was slowest to lift productivity, and why?



PART EIGHT

THE AVERAGE GLOBAL FAMILY, 2000 TO 2026

The new unit: a family not tied to any country, on the eve of the break.

08

What does the world's average family make today, and what has changed since 2000?

The average global family directs about 132 HLE of useful work a year, or 38 per person. Energy per head is flat. Computing power per head has risen ten-thousand-fold, and the first AI work has arrived.

SCOPE

World average

Europe

United States

From 2000 the report changes its unit. Europe was the frontier for the earlier series for two reasons: its data are the richest, and Rome is the starting point. But by 2026 Europe's output per head trails America's. And the biggest recent changes, such as mobile internet, cloud computing and AI, are global. So the rest of the story follows an **average global family**. It uses world averages and is not tied to any single country. It has about 3.5 people. In 2026 it directs about 132 HLE of useful work a year, or about 38 for each person. Almost all of that is done by machines, not by the family itself. This global family is poorer than the Western European family of the earlier chapters. In 2000 the Western European family directed about 290 HLE, but the average global family directed about 125. The 'Average global family' row in the table below uses the global family.

Household

3.5 people

Per person

~38 HLE

Online

74% of people (ITU 2025)

THE WORLD IN A FEW NUMBERS

World output per person rose from about \$12,700 in 2000 to about \$18,800 in 2020 and about \$22,500 in 2026 ^[26]. These are PPP dollars at 2021 prices, meaning dollars adjusted so they buy the same amount in every country. That is a rise of about 1.75 times since 2000. In the 2025 dollars used elsewhere in this report, it comes to about \$26,300 a person in 2026. Average household size fell from about 4.0 people to 3.5, so output per household rose less, by about 1.5 times.

Energy per person hardly moved: it was about 70 gigajoules in 2000 and about 73 in 2025 ^[27]. This report's own estimate for 2026, used in the table below, is about 75. In 2025 the world used more than 600 exajoules of energy in total for the first time. An exajoule is a billion gigajoules. Growth came from getting more output from each unit of energy. Economists call this decoupling, because output and energy no longer rise in step. Electricity per person did rise, by about half, to about 3,780 kWh ^[28]. Internet use went from 7% of people to 74%, or about six billion people ^[29].

METRIC	2000	2010	2020	2026
GDP per head (2021 PPP \$, thousands)	12.7	16.3	18.8	22.5 (est.)
GDP per head (2025 \$, thousands)	15	19	22	26
Primary energy per head (GJ)	70	76	73	75
Electricity per head (kWh)	2,500	3,050	3,400	3,780
Internet users (% of people)	7%	29%	60%	74%
Household size (persons)	4.0	3.8	3.6	3.5
Useful work + delegated thinking per head (HLE)	31	35	36	38
Average global family (HLE)	125	133	128	132

Source: World Bank WDI 2026, IMF WEO April 2026, Energy Institute 2026, Ember 2025, ITU 2025, UN DESA, 311 Institute analysis 2026.

THIRTY-EIGHT SERVANTS EACH


In this report's unit, the average person on Earth in 2026 directs about 37 HLE of useful machine work. It comes from motors, boilers, vehicles and lights, plus about half an HLE of their own labour. In 1950 Fuller counted 38 energy slaves per person on Earth ^[23]. That turns out to be almost exactly today's figure, counted as useful work. The world average has risen only modestly since 2000, by about a quarter from about 31 HLE. Energy per person stayed flat, so most of that

rise came from machines that waste less energy. The average American directs three to four times as many.


The average global person, 2000 to 2026

SOURCE: WORLD BANK 2026, ENERGY INSTITUTE 2026, 311 INSTITUTE ANALYSIS 2026

VALUE, \$ THOUSAND (2025 PRICES)



USEFUL WORK INCL. AI, HLE



THE SERVANTS THAT THINK

What has changed most is computing. Martin Hilbert's work puts the world's general-purpose computing power at around a billion instructions a second per person in 2007. AI accelerators are computer chips built to run AI, and their total computing capacity now doubles about every seven months ^[32]. By the end of 2025 it came to about one to two trillion operations a second for every person alive. Since 2000, computing power per person has risen ten-thousand-fold or more. Most of that rise came after 2020.

Most of that computing is not yet doing work handed over by people. It trains AI models, answers chat questions, ranks social media feeds and renders video. This report finds that in 2026 the average person gets about 0.15–0.5 HLE of accepted thinking work from AI. That is about 400–1,250 hours a year of drafting, coding, searching and summing up. It is spread very unevenly, and only about one working-age person in six uses generative AI at all ^[85]. Generative AI is AI that writes text or code, or makes images. But this is the fastest-growing channel, and it is where the ratio breaks next.

One way to track AI's workload is to count tokens, which are the small chunks of text that an AI model reads or writes. In October 2025, Google said its services handled more than 1.3 quadrillion tokens a month ^[101]. A quadrillion is a thousand trillion. That was up from about 980 trillion in June 2025, though the pace of growth had slowed. Part of the rise came from reasoning models. These are AI models that work through a problem step by step, so they use many more

tokens per question. So tokens measure effort, not accepted work, which is why this report keeps the HLE as its main unit.

Governments are pushing hard. In August 2025, China's State Council, its top government body, issued an 'AI+' plan ^[99]. It targets more than 70% use of AI agents and smart devices by 2027, and more than 90% by 2030. A target is not a result, but it shows how fast the largest economies want this channel to grow.

EUROPE AS A FOLLOWER

Europe needs one more word. On the old measures, such as output per hour and energy per head, Western Europe is still rich. On the new measures, it is a follower. American firms lead in frontier AI models and cloud capacity. Chinese firms lead in humanoid robots, which are robots shaped roughly like a person. IDC, a market research firm, counted about 25,000 humanoids shipped in the first half of 2026, about 78% of them from Chinese makers ^[53]. AI use is highest in the Gulf states and Singapore ^[85]. The European thread served this report well for two thousand years, but for the next seventy-five years the useful unit is the global family. The new frontier is wherever compute (computer processing power), energy and robots are cheapest.

CHAPTER SUMMARY

Since 2000, world output per person has risen about 1.75 times, while energy per person has stayed near 70–75 GJ. Internet use has gone from 7% to 74% of people. AI computing has gone from nothing to over a trillion operations a second per person. The family's own labour is under two HLE. Each person's AI work is about 0.15–0.5 HLE, and it is growing fastest.

The family now directs dozens of energy servants, and the new ones can think.

CHAPTER CLOSE

WHAT NEXT?

The average global family, 2000 to 2026

PREDICTIONS

PREDICTION

75% CONFIDENCE · BY 2030

By 2030 the AI tokens processed per person on Earth reach at least ten times their 2025 level. Tokens are the small chunks of text that an AI model reads or writes.

BASIS

Google's monthly token count rose from about 980 trillion in June 2025 to over 1.3 quadrillion by October. Epoch AI finds AI computing capacity doubles about every seven months, and prices for the same AI ability keep falling. Our analysis expects each person's accepted AI work to grow from about 0.15–0.5 HLE in 2026 to about 0.3–0.6 HLE by 2030. Early sign: large AI providers report token volumes still at least doubling each year in 2027.

TRIPWIRE

Proved wrong if reported AI tokens per person in 2030 are less than ten times their 2025 level.

QUESTIONS TO ASK

- 1 Does your organisation count the machine work it directs, or only its staff?
- 2 How does Europe rank on compute and AI work, the new measures?
- 3 How much of your week is already handed to software that acts alone?



PART NINE

THE COMMAND EQUATION

A simple formula for how much one person can produce when they direct AI agents, robots and machines.

09

If one person no longer sells just their own hours of work, what equation takes the place of the old one?

A person's output becomes their own work plus everything they direct. It is capped by whichever limit hits first: their attention, their budget, their energy or their buyers.

SCOPE

Digital

Physical

Virtual

The old equation of work was so simple that nobody wrote it down. Output equals hours worked times the output of one hour. A person's pay was those hours times a price. That price was set by skill, by how rare the skill was, and by bargaining power. Every job market, tax system and pension scheme in the world is built on this idea. This chapter proposes a new equation to replace it. It is the heart of the report. It sets out what one person can produce once they stop working alone and start to direct AI agents, robots and machines. The chapters that follow test each part of it against the evidence.

First, a reminder of the unit this report uses. One HLE (Human Labour Equivalent) is one adult's working year of output. We count it in one of two ways only. The physical way is about 75 watts of useful work for 2,500 hours, or about 190 kWh. The time way is 2,500 accepted hours of work by a person, a robot or an AI agent. An AI agent is software that can carry out a task on its own, step by step. We never use HLE to say what work is worth, because worth is always shown in dollars.

STATING THE EQUATION

Here is the equation in full. The text below it explains each letter in plain words.

THE COMMAND EQUATION · 311 INSTITUTE

$$O = h \times [1 + \text{sum of } (n \times t \times s \times r) / T]$$

capped by the smallest of four limits: $h \times k$, B / p , E / e and D

O	output a person directs in a year, in HLE
h	the person's own working time (1 for a full-time adult)
sum of	add up over every channel: digital (software, content, services), physical (robots and machines), virtual (VR, AR and spatial worlds)
n, t	units directed in a channel and the hours each runs
s, r	speed relative to a human, and the share of work accepted after review
T	2,500 hours, one HLE
k	direction capacity: HLE of delegated work one person can specify and verify a year
B / p	production budget divided by the price of an HLE-year of delegated work
E / e	energy available divided by the energy per HLE-year
D	demand: what buyers, or the person, will actually absorb

Read the equation from left to right. A person adds their own work, h , measured in HLE. For a full-time adult it is one. On top of that they direct units in three channels. The digital channel holds AI agents and everything they make for ordinary screens, from code and data to songs and videos. The physical channel holds machines and robots. The virtual channel holds work made for VR, AR and spatial worlds, which is tiny today. Each unit runs for some hours and works at some speed compared with a human. Some share of its work is accepted after checking. Divide by the 2,500 hours in one HLE-year and you have that channel's share. Add the channels together and you get what the person could produce if nothing got in the way.

But something always gets in the way, so the right-hand side lists four limits. Output is capped by whichever one bites first. **Direction**: one person can only plan, check and stand behind so much delegated work. This report calls that amount k , or direction capacity. **Budget**: delegated work has a price, so a person can only buy as many HLE as their spending allows. **Energy**: every agent-hour and robot-hour uses power. **Demand**: output that nobody wants is not really output at all.

So the limits can be written very simply. Output can be no bigger than the smallest of four amounts: $h \times k$, B / p , E / e and D . Here B is the budget and p is the price of one HLE. E is the energy on hand and e is the energy each HLE needs. D is what buyers want.

WHY EACH TERM MATTERS

Not one of the letters is random, and each one links to evidence in the chapters that follow. The letter **n** is the number of units. Heavy users of coding agents already run about 5–15 agents at the same time. The letter **t** is hours run. It is limited by how long a task an agent can finish without help. METR, an independent AI testing lab, measures this as a 'time horizon'. That is the length of task an agent finishes half the time. It grew from about two hours in August 2025 to about twelve hours in early 2026 ^[38]. METR first reported 14.5 hours, then rescored it to about 12.

The letter **s** is speed, and it is large. GDPval is OpenAI's test of real expert tasks. On it, models already run about a hundred times faster than the experts, before anyone checks the work ^[39]. The letter **r** is the share of work that is accepted, and it is the weak link. Controlled trials keep showing how much work fails review ^[36].

Now for the limits. The letter **k** grows as agents get more reliable. It also grows as people learn to direct agents that direct other agents. The letter **p** is the price of one HLE-year of digital work, and it falls fast. The price of reaching a fixed level of AI skill falls between 9 and 900 times a year, depending on the task ^[40]. For one common science test it fell about 40 times a year. There is no single typical rate. The letter **e** is the energy per HLE-year, and it falls too. A typical text prompt to Google's Gemini used about 0.24 watt-hours in 2025. That was 33 times less than a year earlier ^[50]. Only **D**, demand, does not clearly rise or fall with technology. It may matter most once the other limits ease.

A WORKED CASE: \$180 A MONTH

The chapter on method, earlier in this report, works through this case step by step. Here is the short version. Take a builder in October 2026 who makes commercial software alone, mainly through AI. They pay about \$180 a month for a top-tier agent plan and keep about ten agents busy for most of the working day. Once the rejected work is thrown away, the agents deliver about **18 HLE** a year. That is the work of about 18 full-time people. It could be up to about 45 if the agents

were faster and more of their work were kept. Hiring those people would cost about \$2.7 million a year.

For this heavy user, budget does not bite, because the plan is a flat fee. Energy does not bite either, and demand does not bite yet for a builder with customers. What bites is **direction**. The person only has so many hours to describe the work, review it and decide what comes next. Ten agents is about as many as one skilled person can keep useful, which matches what other heavy users report. Compute (computer processing power) and energy only start to bite at something like global scale.

THE ORDINARY PERSON'S NUMBERS

The heavy user is not typical. The model in this report follows an ordinary skilled person instead. They spend a tenth of an average income on delegated work, about \$2,600 a year. Four-fifths of it, about \$2,080, buys AI work at about \$120 per HLE. That gives about 17 HLE. The last fifth, about \$520, buys robot work. A robot HLE still costs about \$15,000, so that money buys almost nothing yet, about 0.03 HLE. That is why the AI total is 17, not 22.

Add 1 HLE of their own work and about 37 HLE from their share of everyday machines, such as car motors, power stations and factories. In 2026 that comes to about 56 HLE in all. The heavy user pays the same rate, about \$120 per HLE. But they spend all of their money on AI work, so they get about 18 HLE for less, \$2,160 a year.

For this person the budget runs out first, not attention. In the agentic diffusion branch, budget stays the limit for the ordinary person in every decade to 2100. Attention becomes their limit only in the explosive window, from the 2030s. So today attention limits heavy users, while for most people money does.

WHY CHECKED WORK COSTS MORE

Token prices are falling very fast indeed. A token is a small chunk of text that an AI model reads or writes. But the cost of *accepted* work falls more slowly, for three reasons. First, much AI output is rejected, or has to be checked before anyone can use it. Second, harder tasks use far more tokens than easy ones, so the bill grows as the work gets harder. Third, human checking time does not get cheaper. An hour of a skilled reviewer still costs what it did. So the price of a finished, trusted hour of work falls more slowly than the price list suggests.

Our analysis also found two gaps in what the world measures. The slowest condition for this whole future is direction capacity. That is how much delegated work one person can plan, check and

stand behind. Nobody measures it yet. The energy used per accepted hour of AI work is not measured either. Without these two numbers, leaders are steering by guesswork.

OUTPUT, PRODUCTIVITY AND CAPTURE

The equation measures **output**, the amount of work done. It does not measure **productivity** as economists use the word, which is output per hour worked. For this builder, output per hour has risen about twenty times, which makes the number almost useless. Nor does the equation measure **capture**, which is who keeps the value. A builder who directs 18 HLE might earn more than a team of 18 people. Or the price of their software might fall towards zero, because a thousand others can build the same thing. The later chapters on wealth and work follow the capture question, while these middle chapters follow the amount.

WHAT THE EQUATION LEAVES OUT

Three things need a warning label. First, HLE adds up things that are not alike. An hour of accepted code and an hour of a robot stacking boxes both count as one hour of human work. That is on purpose, because it is what a person could otherwise have hired. But it is a choice, not a law of nature. Second, quality is folded into the share of work accepted. Work that is accepted but wrong will count as output until someone notices.

Third, the equation describes ability, not happiness. Think of a person who directs 30,000 HLE in 2100, as one might in the agentic diffusion branch. They are not 30,000 times happier than a Roman farmer. They can simply cause 30,000 times as much work to happen.

CHAPTER SUMMARY

The command equation says a person's output is their own work plus what their agents and robots do, after checking. It is capped by direction, budget, energy and demand. A heavy user on a \$180-a-month plan directs about 18 HLE, the work of about 18 people. For them, attention is the limit. For most people, money runs out first.

The new unit of work is output directed, not hours sold, and for heavy users its limit is attention.

CHAPTER CLOSE

WHAT NEXT?

The command equation

PREDICTIONS

PREDICTION

65% CONFIDENCE · BY 2030

By 2030 a published survey shows the top tenth of coding-agent users run 25 or more agent sessions at once at peak, up from 5–15 in 2026.

BASIS

Agent time horizons double every 4–7 months, and agents that manage other agents are spreading. Early sign: tools that let one person watch and merge dozens of agent sessions become standard in 2027–2028.

TRIPWIRE

Proved wrong if surveys in 2029–2030 show the top tenth of users still run fewer than 25 sessions at once at peak.

QUESTIONS TO ASK

- 1 How much delegated work can you direct and check in a year?
- 2 Which of the four limits hits your business first?

10

PART TEN

SWARMS OF AGENTS

Digital and thinking work: moving from AI that helps to AI that does the job.

10

How much thinking work can one person direct today, and how fast is that growing?

AI that helps a person gives small, uneven gains. Handing whole tasks to AI agents gives large ones. The length of task agents can finish alone doubles every 4–7 months, while the price of the same skill falls 9–900 times a year.

SCOPE

Evidence

Capability

Cost

The evidence on AI and personal output is messier than fans or critics admit. It becomes clear once you split two very different things. One is AI that *helps* a person do their own work. The other is AI that *does work the person hands over*. Help makes one worker a little faster. Handing over lets one worker act like many, and this chapter looks at both. It uses the best trials, the best tests of what agents can do alone and the real costs of running them. It ends with the small firms that already run on agents, and with the limit that stops one person from directing more.

WHAT THE TRIALS ACTUALLY SHOW

The best-known controlled studies all look at help. In one, an AI assistant was rolled out in stages to 5,179 customer-support agents. It raised problems solved per hour by 14% on average and 34% for beginners. For the best staff it made almost no difference ^[33]. In a randomised trial of 453 professionals, writing tasks took 40% less time and quality rose 18% ^[34].

Consultants at Boston Consulting Group (BCG) tried GPT-4 on their work. On tasks the model was good at, they finished 12.2% more tasks, 25.1% faster and at over 40% higher quality. On tasks outside the model's skill, they were 19 percentage points *less* likely to be right ^[35].

Two other results point the opposite way. In 2025 the research group METR studied experienced open-source developers working in code they knew well. With early-2025 AI tools they were 19% *slower*, while believing they were 20% faster ^[36]. And Danish government records on about 25,000 workers found no measurable effect of AI chatbots on pay or hours ^[37]. Taken together, the evidence says help works best where the task is simple and the worker less expert. It can hurt where the worker is expert and the task is tangled. It has not yet shown up in national accounts.

The table below sums up these trials. In the BCG row, '19 points' means 19 percentage points.

STUDY	SETTING	RESULT
Brynjolfsson, Li & Raymond	5,179 support staff	+14% per hour, +34% for beginners, about 0 for top staff
Noy & Zhang	453 professionals, writing	time -40%, quality +18%
Dell'Acqua et al. (BCG)	758 consultants	tasks AI is good at: +12% tasks, 25% faster, >40% better. Tasks it is not: 19 percentage points less accurate
METR randomised trial	16 expert coders, their own code	19% slower, but believed they were 20% faster
Humlum & Vestergaard	~25,000 Danish workers	no measurable effect on earnings or hours

Source: Quarterly Journal of Economics 2025, Science 2023, HBS/BCG 2023, METR 2025, NBER 2025.

FROM HELP TO HANDING OVER

Handing work over is a different animal. The question is not whether a person types faster. It is how long a task an agent can finish on its own. METR measures this as a 'time horizon'. That is the length of a human task that a model finishes half the time. It roughly doubled every seven months


from 2019 to 2025, and about every four months since ^[38]. So this report takes a doubling time of roughly 4–7 months. The trend line in the chart below doubles every 6 months, in the middle of that range.

GPT-5, in August 2025, managed about 2 hours 17 minutes. Claude Opus 4.5, in December, managed about 4 hours 49 minutes. Claude Opus 4.6, in early 2026, was first reported at 14.5 hours. METR later rescored it at about 12 hours. That is about a day and a half of work per run. METR warns that its task set is close to its ceiling and the error bars are wide. The horizon for tasks finished 80% of the time is about ten times shorter. If METR's task set hits its ceiling, this report's call will be judged on the next set or on another independent tester.

Agent time horizon: how long a task one agent can finish alone

SOURCE: METR 2026 EVALUATIONS, 311 INSTITUTE ANALYSIS 2026

TASK LENGTH FINISHED HALF THE TIME · HOURS · LOG SCALE



A second test points the same way. OpenAI's GDPval asks experts to judge model work against their own, across 44 jobs. The experts have fourteen years of experience on average. In September 2025 the best models won or tied about 48% of the time. By December it was about 71% ^[39]. Before human checking, the models were roughly a hundred times faster and cheaper than the experts. That is the s term in the equation, and it is huge.

WHAT THINKING WORK COSTS

The cost of handing over thinking is falling fast. The price of running a model as good as GPT-3.5 fell more than 280 times between November 2022 and October 2024 ^[41]. Epoch AI finds that the price of reaching a fixed level of skill falls 9–900 times a year, depending on the task. For one test it was about 40 times a year, so there is no single typical rate ^[40].

A top coding agent today uses perhaps one to five million tokens an hour. Most of them are cached, which means stored and reused at a lower price. At 2026 top-tier prices that is about \$3–30 per agent-hour. On mid-tier or open models it is ten cents to a dollar. A US office worker costs \$55–85 an hour once all costs are included.

Flat-fee plans are cheaper still: a plan of \$100–200 a month buys close to full-time use of several top agents. At list prices that use would cost hundreds or even a few thousand dollars in tokens. That is why this report sets the real price for a skilled user at about \$120 per HLE-year in 2026. An HLE-year is 2,500 hours, so an hour of accepted AI work costs the skilled user about 5 cents. A beginner paying list prices could pay \$3,000–30,000 for the same work. The gap is itself a finding, because **skill in directing** agents can make output about 25 to 250 times cheaper.

ONE-PERSON COMPANIES

Tiny teams give the clearest real-world test of the equation. Base44 was started by one founder, Maor Shlomo. It had about eight staff when Wix bought it for \$80 million in cash, about six months after it began ^[42]. Lovable reached about \$400 million of yearly recurring revenue with about 146 staff, TechCrunch reported in March 2026. That is about \$2.7 million per person ^[43]. Cursor passed \$2 billion of yearly revenue in March 2026, Bloomberg reported via TechCrunch ^[44]. Pieter Levels runs a set of products that earn about \$3 million a year on his own ^[45]. A typical listed software company makes \$200,000–400,000 per employee.

The table compares these firms. In the Output column, 'annualised' means a monthly figure multiplied by twelve. The Per head column divides output by headcount.

COMPANY OR PERSON	OUTPUT	HEADCOUNT	PER HEAD
Base44	\$80m cash sale at 6 months	~8 at sale (1 founder)	~\$10m of sale value
Pieter Levels portfolio	~\$3m a year	1	~\$3m
Lovable	~\$400m yearly recurring revenue (Mar 2026)	~146	~\$2.7m

COMPANY OR PERSON	OUTPUT	HEADCOUNT	PER HEAD
Cursor (Anysphere)	\$2bn annualised (Mar 2026)	~300 (2025)	~\$5–10m (estimate)
Typical listed software company	—	—	~\$0.2–0.4m

Source: TechCrunch 2025–26, Bloomberg 2026, Indie Hackers 2025, 311 Institute analysis 2026.

In 2025 Dario Amodei, who runs the AI company Anthropic, predicted a one-person billion-dollar company by 2026. As of October 2026, none has been verified, and one possible case is contested. Medvi, a weight-loss drug marketer, told the New York Times it made about \$401 million in 2025 with about two staff. Those figures are self-reported, not audited, and come from a single profile that critics challenged ^[94].

That matters. It reminds us that output is not value. Selling, trust and timing still decide which delegated output becomes a business. The equation says what one person can make. It does not say anyone will buy it.

Our analysis also suggests that big, existing firms are more likely to absorb agents than to be swept away by a wave of one-person giants. Large firms already have the customers, data and trust. They can add agents to what they do. In the central (agentic diffusion) case, AI agents do most of the world's work-hours by about 2036, while humans still direct and check the work.

THE HUMAN IS THE BOTTLENECK

Heavy users say they run 5–15 agents at once. The limit they name most is their own time to review work. Energy is never one of them. In the equation, k and r bite long before e . This is the most practical finding in these middle chapters. The next decade's gains will come less from smarter models. They will come more from better ways to describe, check and trust delegated work. Think of agents that check agents, tests that run on their own, and track records that let a person accept work without reading every line.

Checking matters most in regulated professions such as law, medicine, audit and engineering. There, a named human must sign off the work, so verification is the gate. Our analysis expects checking and sign-off to be the main thing holding agents back until about 2032. Until then, the number of hours a qualified person can spend checking sets the pace.

CHAPTER SUMMARY

Trials show AI help raises output 14–40% on simple tasks but can slow experts. National figures show no effect yet. The big change comes from handing work over. Agent time horizons rose from about two hours to about twelve in under a year. Tiny teams already earn tens of millions of dollars.

Help makes a person faster, but handing work over makes a person many.

CHAPTER CLOSE

WHAT NEXT?

Swarms of agents

PREDICTIONS

PREDICTION

60% CONFIDENCE • BY 2028

By end-2028 independent testing shows an AI agent finishing 160-hour software or research tasks half the time.

BASIS

METR rescored Opus 4.6 at 12 hours (first 14.5). The 80% horizon is ten times shorter. Early sign: 40 hours in 2027.

TRIPWIRE

Proved wrong if no independent or successor test shows this by end-2028.

PREDICTION

30% CONFIDENCE • BY 2029

By 2029 a company with one full-time human reports \$100m or more of audited yearly revenue.

BASIS

Base44 had about 8 staff at sale. Medvi's claim is unaudited and contested. Early sign: audited solo revenue above \$20m by 2027.

TRIPWIRE

Proved wrong if no such audited case exists by end-2029.

QUESTIONS TO ASK

- 1 Which of your tasks could move from help to handover?
- 2 Where do your people feel faster but measure slower?



PART ELEVEN

FLEETS OF ROBOTS

Physical output: humanoid robots, self-driving machines and the stubborn cost of moving atoms.

When will one person direct a fleet of robots the way they now direct a swarm of AI agents?

Robots already cost less per hour than rich-world wages, but not yet per unit of useful work. Numbers are tiny, supply chains rely on a few countries, and even hopeful forecasts give one humanoid per three people by 2060.

SCOPE

Industrial

Humanoid

Self-driving systems

If agents are the story of the next decade, robots are the story of the next fifty years. The equation treats them the same way. It counts units, hours, speed and acceptance. But the numbers that go in are very different, and so is the limit that bites. Robots must move real things, so they need motors, metals, batteries and safety rules. Each of these takes years to build up, and none of them gets cheaper as fast as computing does. This chapter looks at the robots already at work, the new humanoids, what they really cost, and what bank forecasters expect.

TODAY'S FACTORY ROBOTS

The world installed about 542,000 industrial robots in 2024. It was the fourth year in a row above half a million. That brought the total in use to about 4.66 million ^[52]. China took more than half of the new ones. South Korea has the most robots for its size. In 2023 it had 1,012 robots for every 10,000 factory workers, or about one robot per ten workers ^[52].

Across the whole world, that is roughly one factory robot for every 750 workers. These are fixed machines that do one job, and firms own them. They raise output per factory worker, but they do not yet give ordinary people extra reach.

THE HUMANOID MOMENT

A humanoid robot is a robot shaped roughly like a person, built to do many kinds of task, and these are new. About 13,000–18,000 shipped worldwide in 2025, depending on which tracker you trust. The research firm IDC counts about 25,000 in the first half of 2026 alone, with about 78% made by Chinese firms ^[53]. Another tracker puts China's share above 97%, but it counts a smaller total.

AgiBot built its 10,000th unit in March 2026. Unitree shipped over 5,500 in 2025 and listed its shares on Shanghai's STAR Market. Unitree's G1 sells for about \$16,000, while 1X sells its NEO home humanoid for \$20,000, or \$499 a month. Many hard tasks are still done by a human steering it from far away ^[54]. Tesla's Optimus has slipped towards 2027 for mass production. Hyundai plans to be able to build 30,000 Boston Dynamics robots a year by 2028.

The table below compares what humanoids cost per hour. Capex is money spent up front to buy the equipment, here the robot's price. A bill of materials is the cost of all the parts that go into a machine. 'h/day' means hours a day. 'All-in' means the cost includes upkeep and software, not just the price of the robot.

SCENARIO	CAPEX	LIFE	HOURS A YEAR	ALL-IN \$ PER HOUR
Consumer humanoid (NEO class)	\$20k	5 yr	2,000	~\$3.0
Industrial, two shifts	\$50k	5 yr	5,000	~\$3.5
Industrial, 2030 bill of materials	\$15k	7 yr	6,000	~\$0.6
Western premium, today	\$100k	5 yr	4,000	~\$8.8
Lease at \$499 a month, 12 h/day	—	—	4,380	~\$1.4

SCENARIO	CAPEX	LIFE	HOURS A YEAR	ALL-IN \$ PER HOUR
US median wage (comparison)	—	—	—	~\$23

Source: 1X 2025, Goldman Sachs 2024, Figure AI 2025, BLS, 311 Institute analysis 2026 (our estimate, maintenance and software 10–15% a year included).

CHEAPER PER HOUR, DEARER PER TASK

Do the sums and humanoids already look cheap. Take a \$20,000 home robot that lasts five years and works 2,000 hours a year, and it costs about \$3 an hour. A factory robot working two shifts costs about \$3.50. A premium Western machine costs about \$9, and electricity adds only five to ten cents an hour. A lease at \$499 a month works out at about \$1.40 an hour if the robot runs twelve hours a day. Typical US pay is about \$23 an hour.

Yet humanoids are not yet cheaper *per unit of useful work*. In trials they often work at 30–50% of human speed. They need a human to step in every few minutes or every few hours. They stop for battery cycles of two to five hours. In the equation, s and r are both well below one. This chapter's last prediction asks when that changes in a US warehouse or factory. It would be more likely, about 65%, if a job in any OECD country counted. The OECD is a club of 38 mostly rich countries.

Our analysis looks at the slowest step, because the slowest step decides. It expects humanoids that can work without help around 2031. It puts cost parity, the point where a robot costs the same as a human worker, around 2035. But that is close to a coin toss, as this chapter's last prediction shows. Our analysis expects robots that people accept in their homes around 2038.


WHAT THE FORECASTERS EXPECT

Forecasts about humanoids are wide and mostly hopeful. Goldman Sachs expects a \$38 billion humanoid market by 2035 ^[58]. The Swiss bank UBS expects about 300 million humanoids in use by 2050 ^[57]. Morgan Stanley, in May 2025, forecast a \$5 trillion market by 2050 ^[56]. Bank of America expects about three billion humanoids by 2060. It says more people will own a humanoid than a car. It expects each one to cost \$13,000–17,000 by the early 2030s ^[55]. Divide by the number of people and even this hopeful case gives about one humanoid for every three people by 2060.


Humanoids per 100 people: forecasts and this report's branches

SOURCE: UBS 2025, MORGAN STANLEY 2025, BANK OF AMERICA INSTITUTE 2026, 311 INSTITUTE ANALYSIS 2026

BANK AND ANALYST FORECASTS



THIS REPORT: HUMANOID PER 100 PEOPLE, FOUR BRANCHES



MAGNETS, MOTORS AND SAFETY RULES

Physical output depends on physical supply chains. A robot's joints are moved by small motors called actuators. Many of these use strong magnets made from rare earths, a group of hard-to-process metals. China makes about 94% of the strong permanent magnets that robot motors need ^[59]. Its export limits in April 2025 hit Tesla's Optimus plans directly. Wider limits were announced in October 2025. A US–China deal paused them for one year, and that pause runs out around November 2026.

Refining other key minerals is also packed into a few places. The top three refining countries handle about 86% of it ^[92]. Safety rules lag behind as well, because ISO 25785-1, the first standard for legged factory robots, is still an early draft ^[60]. Not one of these limits holds back an AI agent. All of them hold back a robot.

RENTING THE PHYSICAL WORLD

For one person, the strongest physical helpers in 2026 are systems they do not own. Waymo said in March 2026 that it was giving about 500,000 paid robotaxi rides a week. It aims for a million a week by the end of the year ^[61]. Xiaomi opened a 'lights-out' factory in Beijing in July 2024. It runs with no people on the floor and can make about one phone a second at full speed ^[62]. In 2022 John Deere set a goal of a fully self-driving corn and soybean system by 2030 ^[63].

A person who can book robot-hours, factory-hours and drone-hours by the job directs far more physical output than one who owns a single humanoid. That is the shape of physical command to 2100: **rented fleets, not owned servants.**

After 2050 this report does not give one forecast. It gives four branches with no odds attached. In the agentic diffusion branch, general-purpose robots reach about one per person by 2100. Together they add about 4.2 HLE per person. Add the much larger stock of ordinary machines, the 'energy servants' from the chapter on energy servants, which grow with energy use. Physical output per average person then reaches about 190 HLE, five times today's.

The other branches give less than this. In the setback branch, where a hard shock in the 2040s costs a decade, physical output is about 49 HLE per person. In long drift it is about 69. In the explosive window it is about 450, with five robots per person. Any of these is a big rise for atoms, but next to the digital channel it is small.

CHAPTER SUMMARY

About 4.7 million robots work in factories. IDC counted about 25,000 humanoids shipped in early 2026, mostly Chinese. A humanoid costs about \$1.40–9 an hour, but it is slower and needs help, so useful work still costs more. Banks forecast 3–30 humanoids per hundred people by 2050–2060. Most people will rent robots, not own them.

Bits grow with energy and ideas, but atoms grow with magnets, motors and minerals.

CHAPTER CLOSE

WHAT NEXT?

Fleets of robots

PREDICTIONS

PREDICTION

30% CONFIDENCE • BY 2032

World humanoid shipments pass 1 million units in a single year by 2032.

BASIS

About 25,000 shipped in early 2026. Makers over-promise, so we lowered our first estimate of 50%. Early sign: 250,000 in 2029.

TRIPWIRE

Proved wrong if no year to 2032 passes 1 million units shipped.

PREDICTION

45% CONFIDENCE • BY 2035

By 2035 a leased humanoid undercuts median US workers per accepted unit in one warehouse or factory task.

BASIS

Humanoids cost \$1.40–9 an hour but work slower and need help. Early sign: a US pilot reaches cost parity by 2033.

TRIPWIRE

Proved wrong if no US task passes this cost test by end-2035.

QUESTIONS TO ASK

- 1 Which physical tasks could you rent as robot-hours soon?
- 2 How much do your machines rely on one country's magnets?

12

PART TWELVE

INFINITE WORLDS

Virtual output: VR, AR, smart glasses, spatial worlds and digital twins.

12

What happens to output when whole 3D worlds cost almost nothing to make, and who will step inside them?

Virtual output is anything made for VR, AR and spatial worlds. It is still tiny, about 1% of AI work in 2026.

Smart glasses are selling fast, headsets are not, and AI can now build a 3D world from a sentence.

SCOPE

Glasses and headsets**Generated worlds****Digital twins**

The third channel is the newest and the smallest. Virtual output means anything made in or for virtual reality (VR), augmented reality (AR), mixed reality and spatial computing. VR puts a person inside a computer-made world through a headset. AR lays computer images over the real world, often through smart glasses. Spatial computing means apps and worlds you move around in, not just look at. A song, a video or an article made by AI is digital output, and so is a game on a phone or a TV. Virtual output starts only when the work is built for a headset, glasses or a 3D space. By that test it is about 1% of all AI work in 2026.

A SMALL MARKET TODAY

Virtual output needs a screen on your face, so start with the devices. The research firm IDC counts headsets and smart glasses together as XR (extended reality) devices. XR shipments grew about 44% in 2025, driven mainly by smart glasses. Meta, which owns Facebook, held about 72% of the market, yet shipments of its Quest headsets fell about 42% ^[106].

IDC counted about 2.25 million smart glasses shipped in the first three months of 2026, up 167% on a year earlier ^[107]. For 2026 it expects about 13.6 million smart glasses and 3.2 million mixed reality headsets. By 2030 it expects about 27.3 million glasses and 10.4 million headsets ^[108]. Phones sell more than a billion a year, so this is still a small corner of the market.

GLASSES ARE WINNING, HEADSETS ARE NOT

EssilorLuxottica, which makes Ray-Ban, sold more than 7 million Meta smart glasses in 2025. It had sold about 2 million in all from late 2023 to early 2025 ^[110]. Counterpoint Research found AI glasses shipments rose 263% in the first half of 2026. But about 96% of them had no display ^[109]. They can hear, see and talk, but they cannot yet lay images over the world. So most of today's glasses are a doorway to AR rather than AR itself.

Building this market has been very costly. Meta's Reality Labs, which makes its headsets and glasses, lost about \$19.2 billion in 2025 on revenue of about \$2.2 billion ^[111]. That is almost nine dollars lost for every dollar of sales. In the second quarter of 2026 it lost about \$4.6 billion on \$431 million of revenue ^[112]. One IDC estimate says Apple shipped only about 45,000 Vision Pro headsets in the last three months of 2025, against about 390,000 in 2024 ^[113]. That is a single, disputed estimate, so treat it with care. Samsung's Galaxy XR headset launched in October 2025 at \$1,800 ^[114].

WORLDS MADE BY AI

The making side is moving faster than the devices. In January 2026 Google opened Project Genie, built on Genie 3, a 'world model' that turns a prompt into a world you can move through. Each world lasts about 60 seconds for now ^[118]. World Labs' Marble, launched in November 2025, builds 3D worlds that can be viewed in VR on Vision Pro and Quest 3 ^[119]. Roblox's Cube 3D, shared openly in March 2025, makes 3D objects from a short text prompt ^[120].

Roblox shows what 3D creators can earn. It paid creators about \$477 million in the last three months of 2025, up 70% ^[116]. Its creators earned more than \$1 billion in the year to March

2025 ^[117]. But Roblox is played mainly on flat screens, not in headsets. By this report's definition that makes it digital output, not virtual. The tools to make 3D worlds are arriving before most people own a device to enter them.

DIGITAL TWINS AT WORK

Not all virtual output is for fun. BMW's Virtual Factory, built on Nvidia's Omniverse platform, holds digital twins of more than 30 of its production sites ^[115]. Planners walk through a new line in 3D before anyone moves a machine. BMW expects this to cut planning costs by up to 30%. Training worlds where robots practise before entering a real kitchen are close cousins. GDP, the value of everything a country makes in a year, often misses this work. It shows up later, as a cheaper factory or a better robot.

GENERATED MEDIA COUNTS AS DIGITAL

The flood of AI-made media is real, but it is digital output. About 99,000 new tracks reach streaming services each day, says the data firm Luminate ^[65]. The AI music service Suno is reported to make a Spotify catalogue's worth of songs, roughly 100 million tracks, every two weeks ^[64]. That is about seven million a day, seventy times what all services receive.

Music: AI songs made against all new uploads, per day

SOURCE: LUMINATE 2025, DJ MAG 2025, 311 INSTITUTE ANALYSIS 2026

thousand tracks a day

8,000

6,000

4,000

2,000

0

7,000

99

All new tracks on streaming

Made on one AI app

Making is not listening. Deezer says fully AI-made tracks passed half of its daily uploads in June 2026, yet they are only about 1–3% of streams. Up to 85% of the streams on AI tracks were flagged as fraud, meaning fake plays by bots ^[95]. A study by the firm Graphite, reported by Axios, found that about half of new web articles are written by AI. The share has levelled off, and spotting AI text is not exact ^[66]. Supply has exploded, but real attention has barely moved. The same will happen to 3D worlds.

I ATTENTION, TRUST AND PROOF

Games earned about \$189 billion in 2025, mostly on ordinary screens ^[67]. The creator economy is the money made by people who post content online. Goldman Sachs expects it to grow from about \$250 billion in 2023 to about \$480 billion by 2027 ^[68]. These markets grow with income and free time, not supply. People have about sixteen waking hours a day, and the population peaks at about 10.3 billion in the mid-2080s ^[86].

So three things become rare: **attention**, **proof of origin** and **identity**. Google's Pixel 10 phone signs every photo with C2PA Content Credentials, a digital label that records where an image came from ^[69]. In virtual worlds a fourth limit is added, **comfort**. People must be willing to wear the device for hours.

I HOW BIG VIRTUAL COULD GET

This report's model puts about 1% of a skilled person's production budget on virtual work today. In agentic diffusion that rises to about 4% by 2040 and a fifth by 2100. In 2040 on the central path, a skilled person directs about 11 HLE of virtual work, about 4% of their 280 HLE. Digital work is about 74%, and machines and their own work the other 22%.

After 2050 we give branches with no odds. By 2100 in agentic diffusion, virtual output is about 700 HLE per average person, about 19% of their total. A skilled person directs about 7,500 HLE of it, about a quarter of all they direct. Per average person it is about 0.7 HLE in setback, 2.8 in long drift and 32,000 in the explosive window. Demand bites hardest here, because virtual work only counts if someone puts on the glasses.

I CHAPTER SUMMARY

Virtual output is work made for VR, AR and spatial worlds, and it is about 1% of AI work today. Smart glasses sales tripled in 2025 while headset sales fell. AI can now build 3D worlds from a prompt, and factories plan in digital twins. The flood of AI songs and articles is digital output, not virtual. In agentic diffusion, virtual work reaches about a fifth of output by 2100.

The rare thing is not a world to step into, but a reason to put the glasses on.

CHAPTER CLOSE

WHAT NEXT?

Infinite worlds

PREDICTIONS

PREDICTION

30% CONFIDENCE · BY 2030

By 2030 smart-glasses and XR headset shipments together pass 50 million units in a single year.

BASIS

IDC expects about 16.8 million in 2026 and 37.7 million in 2030, so this needs a third more than IDC's own forecast. Glasses have beaten forecasts lately, but headset sales are falling. Early sign: combined shipments above 25 million in 2027.

TRIPWIRE

Proved wrong if no year from 2026 to 2030 sees 50 million glasses and headsets shipped, counted as IDC counts them today.

QUESTIONS TO ASK

- 1 Which of your sites or products would gain from a digital twin?
- 2 What would make your customers wear smart glasses all day?
- 3 Can you prove what is yours once copying costs nothing?

13

PART THIRTEEN

WHAT BINDS

The real limits on what one person can produce: direction, budget, energy and demand.

13

If making things is no longer the limit, what is?

For a skilled heavy user, the limit is how much work they can direct. For most people it is money. Energy, chips and minerals only limit whole societies, and in the end, demand limits everything.

SCOPE

Individual

Society

Planet

A heavy user can build commercial software for \$180 a month. Such a person is not held back by compute or energy until something like global scale. This chapter tests that claim against all four limits in the command equation, this report's sum for how much one person can direct. The four limits are direction, budget, energy and demand. It tests them at three scales: the person, the society and the planet. Each limit bites at a different scale and at a different time. Knowing which one bites first tells a leader where to spend effort, and where spending is wasted.

DIRECTION: THE HUMAN LIMIT

For the skilled person today, the limit that bites is direction. That means deciding what to build, checking what comes back and taking responsibility for it. METR, an independent AI testing lab, found that expert coders thought AI sped them up when it slowed them down ^[36]. A study of consultants at Boston Consulting Group showed how quality falls apart on tasks a model, the AI system itself, does badly ^[35]. Both of these are failures of direction.

This report sets one person's direction capacity, k , at about 20 HLE in 2026. On the central path it reaches about 230 by 2040. By 2100 it is about 1,000 in the setback branch, 3,000 in long drift, 110,000 in agentic diffusion and 2 million in the explosive window. It grows as agents become more reliable and start to supervise other agents. It also grows as track records let people accept work without reading it. That growth is an assumption, and it is the most important one in the model.

Our analysis finds that direction capacity is the slowest of all the conditions this future needs, yet nobody measures it. No statistics office, employer or survey tracks how much delegated work one person can plan, check and stand behind. The energy used per accepted hour of AI work is not measured either. Both gaps can be closed with simple surveys and logs.

BUDGET: THE ORDINARY LIMIT


For an ordinary person, money runs out first. Say they spend a tenth of an average income on production. In 2026 that tenth is about \$2,600 a year. A skilled user pays about \$120 for one HLE-year of digital work. Four-fifths of the budget, about \$2,080, goes on AI work and buys about 17 HLE. The last fifth goes on robot work, which still costs about \$15,000 per HLE, so it buys almost nothing yet. Add their own work and their share of everyday machines, and they get about 56 HLE in all. The price falls in every branch. By 2100 it is about \$20 in the setback branch, \$10 in long drift, \$1 in agentic diffusion and ten cents in the explosive window. At \$1, it is close to the floor set by energy and hardware.

Income rises too, so in the agentic diffusion branch, budget stays the main limit for the ordinary person in every decade to 2100. Direction becomes the limit only in the explosive window, from the 2030s, and for heavy users today. That is why *who can afford to direct output* becomes a big political question in the chapters on wealth and work.

Which limit bites first for the ordinary person, agentic diffusion branch

SOURCE: 311 INSTITUTE ANALYSIS 2026 (ESTIMATE)

BUDGET LIMIT AGAINST DIRECTION LIMIT • HLE • LOG SCALE



ENERGY, CHIPS AND MINERALS

For whole societies, physical limits are real. Data centres are large buildings full of computers that run AI and online services. They used about 415 TWh (terawatt-hours) of electricity in 2024. That was about 1.5% of the world's total. The International Energy Agency (IEA) expects about 945 TWh by 2030 and about 1,200 TWh by 2035 in its base case. That is roughly 3% of world demand ^[51].

Big tech firms planned capital spending, or capex, that topped \$700 billion in 2026 ^[70]. Much of it goes on chips and data centres. Right now the tightest point is advanced packaging, the step that joins chips together, not chip design. Taiwan's TSMC reportedly plans to double this capacity from about 130,000 to 260,000 wafers a month by the end of 2028 ^[91]. Mineral refining sits in a few places. The top three refining countries handle about 86% of key minerals ^[92].

These limits shape national plans and company profits. They do not stop one person running ten agents, but they do reach ordinary people through power bills. In one auction run by PJM, which manages the grid for much of the eastern US, the price for power capacity rose about ninefold. PJM blamed mainly closing power plants and rule changes. Its market monitor said data centres added about \$9.4 billion to that auction's cost ^[97]. As demand grows, more people will feel this.

THE PLANET'S ENERGY LIMIT

At the scale of the planet, energy sets the outer edge. Humanity used over 600 exajoules of energy in 2025, about 73 gigajoules per person ^[27]. An exajoule is a billion billion joules, and a gigajoule is a billion joules. That total works out at about 19 terawatts of power used all the time. A kilowatt is about the power of a microwave oven, and a terawatt is a billion kilowatts. Sunlight delivers about 85,000–90,000 terawatts to the ground.

Suppose energy use grows 2% a year to 2100, so that it reaches about 80 terawatts. That is only 0.1% of the sunlight reaching the ground. It is about 7.8 kilowatts per person, roughly three times today. That fits the digital output of the agentic diffusion branch easily. In that branch, a kilowatt of data-centre power per person delivers thousands of HLE a year.

Growth of 5% a year would be different. It would need about 1% of the sunlight reaching the ground and would run into limits of land, heat and materials. Over centuries, waste heat would cap any growth that needs more energy, as the physicist Tom Murphy showed in 2011 ^[90]. So very long-run growth must come from using energy better, not from using more of it.

ENVELOPE	VALUE	SHARE OF SURFACE SUNLIGHT
Human primary energy, 2025	over 600 EJ (~19 TW)	~0.02%
2100 at 2% a year	~2,600 EJ (~80 TW, ~7.8 kW per person)	~0.1%
2100 at 5% a year	~23,000 EJ (~730 TW)	~1%
Data centres, 2024 / 2030 / 2035 base	~415 / ~945 / ~1,200 TWh	—
Data-centre power per person, 2030	~13 W	—
Compute per person, 2100, four branches	60 W / 100 W / 1 kW / 3 kW	all fit the 2% envelope

Source: Energy Institute 2026, IEA 2025, 311 Institute analysis 2026 (our estimate).

DEMAND: THE FIXED LIMIT

The one limit that technology does not loosen is demand, because people have only sixteen waking hours. The population peaks at about 10.3 billion in the mid-2080s ^[86]. Much valuable work, such as care, trust, judgement and attention, is valuable because a person gives it.

The equation says a single person in 2100 could direct tens of thousands of HLE in the agentic diffusion branch. Whether anyone wants that much of their output is a market question. Most of it

will be used by the person themselves, by the firms and machines they serve, or by other machines. That is not a reason to doubt the amount. It is a reason to expect its price to fall.

CHAPTER SUMMARY

Direction capacity limits the skilled person today. In agentic diffusion, budget limits the ordinary person to 2100. Energy, chips and minerals limit nations and the planet. Data centres used about 1.5% of world electricity in 2024 and may use about 3% by 2030–2035. Energy growing 2% a year reaches about 80 terawatts by 2100, a tiny share of sunlight. Demand, attention and trust are the ceiling that does not move.

Compute and energy are not your limit, but your attention and your buyers are.

CHAPTER CLOSE

WHAT NEXT?

What binds

PREDICTIONS

PREDICTION

20% CONFIDENCE · BY 2035

Data centres use 5% or more of world electricity by 2035.

BASIS

The IEA base case is about 945 TWh in 2030 and about 1,200 TWh in 2035, roughly 3% of world demand. Big tech spending topped \$700bn in 2026. Early sign: the IEA raises its 2035 base case above 1,600 TWh before 2030. Early warning: a share under 3.5% in 2032.

TRIPWIRE

Proved wrong if the IEA reports data centres below 5% of world electricity in 2035.

QUESTIONS TO ASK

- 1 Which limit holds your organisation back now, and which will in 2035?
- 2 Are you raising direction capacity as carefully as you buy compute?
- 3 How exposed is your AI plan to national energy and supply limits?

14

PART FOURTEEN

ONE PERSON IN 2100

Four branches, decade by decade, and the headline number.

14

In 2100, how much output can one person produce across digital, physical and virtual work?

In 2100 one skilled person could direct about 260 HLE (setback), 700 (long drift), 30,000 (agentic diffusion) or 2 million (explosive window) a year. One of Rome's richest freedmen directed about 33,000 with 4,116 slaves and 3,600 pairs of oxen.

SCOPE

Setback**Long drift****Agentic diffusion****Explosive window**

This chapter brings the whole model together. It does not give one forecast, because the uncertainty is itself the finding, so it uses scenarios instead. Each one is a consistent set of assumptions about growth, energy, robots, compute, prices and direction capacity. The chapter moves decade by decade, first for the average person and then for one skilled person. It ends with the headline number, what moves it, and what decides which branch the world takes. The full values are in the data book at the end of the report. Treat every number after 2050 as one branch among four, not as a forecast.

Setback

~262 HLE

Long drift

~705 HLE

Agentic diffusion

~30k HLE

Explosive window

~2.0m HLE

Up to 2050 we treat three of them as a band. Long drift is the low edge, agentic diffusion the central case and the explosive window the high edge. We judge the real outcome about as likely to fall below the low edge as above the high edge, about one chance in ten each. After 2050 the band ends, and from then on there are four **branches**, and we give no odds for any of them. The fourth branch is the **setback**. In it, a hard shock in the 2040s costs a decade, and recovery after it is slow. The shock might be a major war, a pandemic or a grid failure.

FOUR BRANCHES TO 2100

The table below sets out the main settings of each scenario side by side.

ASSUMPTION	SETBACK	LONG DRIFT	AGENTIC DIFFUSION	EXPLOSIVE WINDOW
Growth in income per head	Lost 2040s decade, then 1.5% a year	1.5% a year	2.5% to 4.5%, then 3%	6% to 12%, then 4%
Energy per head, 2100	~85 GJ	~108 GJ	~264 GJ	~554 GJ
Helper robots per person, 2100	0.05	0.1	1	5
Computer power per person, 2100	60 W	100 W	1 kW	3 kW
Energy per accepted hour of AI work, 2100	8 Wh	5 Wh	0.5 Wh	0.05 Wh
Price of one HLE of digital work, 2100	\$20	\$10	\$1	\$0.10
Direction capacity k, 2100	~1,000	~3,000	~110,000	~2,000,000
Income per head, 2100 (2025 \$)	~\$53k	~\$79k	~\$373k	~\$3.0m

Source: 311 Institute analysis 2026 (our estimate), anchors from UN WPP 2024, IEA 2025, Davidson 2021, Epoch AI 2025, bank humanoid forecasts 2024–26.

Setback follows the long drift until the 2040s. Then a shock wipes out about ten years of progress, and growth restarts at about 1.5% a year per person. **Long drift** extends the twentieth century's world growth of about 1.5% a year per person. Energy per person rises slowly, and robots reach about one per ten people. The price of delegated work falls only tenfold from today's skilled-user level. Direction capacity grows slowly because reliability stalls and institutions resist.

Agentic diffusion sees growth speed up to 2.5–4.5% a year as agents and robots spread, then settle at about 3%. Energy grows about 2% a year, robots reach one per person, and prices fall close to the energy floor. **Explosive window** follows the full-automation studies of Tom Davidson and of Epoch AI ^[87] ^[88]. Epoch's model finds explosive growth once 50–70% of tasks are automated. In this branch,

growth runs at 6–12% a year for decades, then settles near 4%. Energy grows about 3% a year, there are five robots per person, and agents manage other agents at huge scale.

THE AVERAGE PERSON, DECADE BY DECADE

The first result is output per average person, across all channels. In 2026 it is about 38 HLE. Most of that is machine work. It also includes about half an HLE of the person's own labour and a small slice of delegated thinking. On the central path it passes 100 HLE in the late 2040s.


After 2050 the branches split apart quickly. In agentic diffusion, output per person passes 500 around 2070 and reaches about **3,700 HLE by 2100**. It is about 62 in the setback branch, about 100 in long drift and about 106,000 in the explosive window. Households shrink to about three people, so an average family in 2100 in agentic diffusion makes about 11,000 HLE. That compares with about 130 for an average family worldwide today, and about 10 for the Roman family of four. The family column in the table below is this average global family of about three and a half people. It is not the Western European family of four from the history chapters, which made about 294 HLE in 2000.

Output per average person, 2000 to 2100

SOURCE: WORLD BANK 2026, ENERGY INSTITUTE 2026, 311 INSTITUTE ANALYSIS 2026 (ESTIMATE)

HLE PER PERSON · LOG SCALE · BAND TO 2050, THEN FOUR BRANCHES

HLE per person



YEAR	SETBACK	LONG DRIFT	AGENTIC DIFFUSION	EXPLOSIVE WINDOW	AVERAGE GLOBAL FAMILY (DIFFUSION)
2000	31	31	31	31	125
2010	35	35	35	35	133
2020	36	36	36	36	128
2026	38	38	38	38	132
2030	39	39	44	53	151
2040	41	44	65	118	216
2050	38	50	126	1.1k	403
2060	41	56	241	3.3k	747
2070	44	63	548	11k	1.7k
2080	49	73	1.0k	23k	3.0k
2090	54	86	1.9k	49k	5.7k
2100	62	105	3.7k	106k	11k

Output per average person, HLE a year. To 2050 long drift to explosive window is our low-to-high band (we put about a one-in-ten chance below the low line and above the high line). After 2050 the four columns are branches with no odds. Source: 311 Institute analysis 2026 (observed to 2026, our estimates after).


OUTPUT MOVES TOWARDS BITS

Today almost all of a person's output is physical: machines burning fuel. By 2100 in the agentic diffusion branch, digital work is about 76% of what the average person gets in total, machines included. Virtual output is about 19% and physical work, with the person's own, about 5%. Physical output still rises about fivefold, but it is simply overtaken. Atoms stay expensive, and bits do not. What one skilled person directs has a different mix, set out below.

Where output comes from: share by channel, agentic diffusion branch

SOURCE: 311 INSTITUTE ANALYSIS 2026 (ESTIMATE)

SHARE OF HLE PER PERSON



THE HEADLINE: ONE PERSON IN 2100

The second result is the one this report set out to find. What can **one person** produce if they choose to? The model's ordinary producer is a skilled person. They spend about a tenth of an average income on delegated work. A fifth goes on physical work. The virtual share, about 1% of the budget today, goes on VR, AR and spatial work, and the rest on digital work. They also direct their share of the world's machines. In 2026 that person gets about 56 HLE. That is 1 of their own work, about 17 of AI work, almost nothing yet from robots and about 37 from their share of machines. The heavy user in the \$180-a-month case directs about 18 HLE of AI work alone. The next page shows what the same skilled person could direct in 2100 in each of the four branches.

THE NUMBERS

One person in 2100

262 HLE

SKILLED PERSON, SETBACK

311 INSTITUTE 2026 · FORESIGHT

705 HLE

SKILLED PERSON, LONG DRIFT

311 INSTITUTE 2026 · FORESIGHT

30k HLE

**SKILLED PERSON, AGENTIC
DIFFUSION**

311 INSTITUTE 2026 · FORESIGHT

2.0m HLE

**SKILLED PERSON, EXPLOSIVE
WINDOW**

311 INSTITUTE 2026 · FORESIGHT

3.7k HLE

**AVERAGE PERSON, AGENTIC
DIFFUSION**

311 INSTITUTE 2026 · FORESIGHT

33k HLE

ISIDORUS, ROME, 8 BC

PLINY THE ELDER · SINGLE-SOURCE


In the agentic diffusion branch, one ordinary skilled person in 2100 can produce about **30,000 HLE a year**. That is roughly 22,000 HLE of digital work and 7,500 of virtual output. It also includes about 200 of physical work. So what this person directs is about 74% digital, 25% virtual and about 1% physical. About 15 of that comes from robots they rent, and the rest from their share of ordinary machines.

That equals the unaided labour of thirty thousand adults. Once machines are counted, it is roughly what about 800 people get today, at about 38 HLE each. It is also the work of about three thousand Roman families of four. It is close to what the Roman Gaius Caecilius Isidorus directed through 4,116 slaves and 3,600 pairs of oxen, about 33,000 HLE ^[5]. The difference is that the person in 2100 owns nobody. They do it with about a tenth of an average income.

What one person can direct, 2100, by branch

SOURCE: 311 INSTITUTE ANALYSIS 2026 (ESTIMATE)

LOG SCALE, BRANCHES WITHOUT ODDS



What limits that number depends on the branch. In the setback, long drift and agentic diffusion branches, **budget** bites. The person could direct more if delegated work were cheaper or they spent more. In the explosive window, **direction** bites instead. Delegated work is almost free, and the limit is how much one person can plan and stand behind. That is about 2 million HLE, the unaided labour of two million adults.

Some people have money to spare, so budget never limits them, and we call this a capital-backed frontier person. By 2100 they reach about 1,000 HLE in setback, 3,100 in long drift, 110,000 in agentic diffusion and 2 million in the explosive window. Energy does not limit any single person in any branch. Demand limits most kinds of output in every branch. That is why measured value rises far less than volume.

VOLUME RISES FAR FASTER THAN VALUE


The third result is the gap between how much is made and what it is worth. In the agentic diffusion branch, output per person rises about 100 times between 2026 and 2100. Measured income per person, in 2025 US dollars adjusted for prices in each country (PPP), rises about 14 times. It goes from about \$26,000 to about \$370,000.

So measured value understates output volume by about 1.7 times by 2060 and about 7 times by 2100. The difference comes from falling prices, the same fall that made Nordhaus's 'lumen-hour' of light almost free ^[22]. This time it applies to thinking, content and, later, much physical work. This gap is the bridge to the chapters on who gains. When output is plentiful, wealth goes to whatever stays rare.

Volume against value, agentic diffusion (2026 = 1)

SOURCE: 311 INSTITUTE ANALYSIS 2026 (ESTIMATE)

INDEX, LOG SCALE



WHAT MOVES THE NUMBER

Three assumptions move the headline number the most. **The price of delegated work in 2100:** halve it and the agentic diffusion figure roughly doubles, until direction starts to bite. **Direction capacity:** suppose agents never learn to manage other agents reliably, and k stalls near 3,000. Then the explosive window falls back towards long drift. **Budget share:** a person who spends a fifth of their income on production directs twice as much. Energy and robot numbers matter much less for the headline, because the digital channel is so large. These branch figures are kept open on purpose. They will need a great deal of refining as evidence arrives.

WHAT DECIDES THE BRANCH

No one can say today which branch the world will take. But we can name the branch points: the few things that decide it. For each one we give checkpoints for 2030 and 2035. These let a reader see which branch we are moving towards.

Direction capacity. Can people learn to direct far more work, and can agents be trusted to check each other? By 2030, watch whether top users run 25 or more agents at once. By 2035, watch whether top users run about 50 or more agent sessions at once, roughly what agentic diffusion implies. Our analysis then puts direction capacity at about 39 HLE in long drift, 96 in agentic diffusion and 180 in the explosive window. Today it is about 20.

The price of accepted work. By 2030, an HLE-year of digital work costs a skilled user \$98 in long drift, \$66 in agentic diffusion and \$45 in the explosive window. By 2035 those prices are about \$77, \$32 and \$13.

Energy. Can grids keep up with data centres? By 2030, watch whether data centres stay near the IEA's path of about 945 TWh. By 2035, watch whether power prices near data centres keep rising, as they did in PJM's auction.

Robot cost. By 2030, watch whether humanoid shipments reach the hundreds of thousands a year. By 2035, watch whether a leased humanoid beats a US worker on cost for at least one real task.

Who owns compute. Open-weight models are AI models whose inner settings anyone can download and run. Tokens are the small chunks of text an AI reads or writes. OpenRouter, a site that sends requests to many AI models, found open-weight models handled about a third of its tokens in December 2025 ^[73]. By 2030, watch whether that share holds. By 2035, watch whether a handful of firms control most of the world's AI computing.

A 2040s shock. This cannot be forecast, but its risk can be lowered. In 2030 and 2035, watch spare grid capacity, pandemic readiness and how tense relations are between the great powers.

DECISIONS THAT HOLD IN EVERY BRANCH

Some choices pay off whichever branch the world takes. Build the skill of directing agents: describing work clearly, checking it and standing behind it. Measure delegated output and accepted work, not just hours. Keep records that show which work was checked, so trust can grow. Keep hiring and training juniors, because tomorrow's experts start as today's beginners. Avoid relying on one supplier for chips, magnets or computing. Plan for a world where output is cheap but trust is rare. Each of these is useful in the setback branch and in the explosive window alike.

ONE SHARED BET

Our analysis found that most of this report's calls are really one shared bet. The bet is that AI agents become reliable enough to be trusted without line-by-line checking. If that holds, many calls come true together, and if it fails, many fail together. Readers should treat the calls as linked, not as separate coin tosses.

WHAT WE CAN RULE OUT

Some outcomes can be ruled out with confidence. One ordinary person will not direct a billion HLE by 2100, because energy and physics do not allow it. There will not be 100 humanoids per person by 2060. Physical output will not grow as fast as digital output. What is left is still wide, from a few hundred HLE to a few million, and that width is the honest finding.

CHAPTER SUMMARY

By 2100 a skilled person could direct about 260, 700, 30,000 or 2 million HLE a year, across four branches with no odds. In agentic diffusion, what the skilled person directs is about 74% digital, 25% virtual and about 1% physical. The average person makes about 3,700 HLE, against about 10 for a Roman family of four. Volume per person rises about 100 times, income about 14 times.

One person in 2100 could direct the unaided work of thousands, so the question becomes who decides what is made.

CHAPTER CLOSE

WHAT NEXT?

One person in 2100

PREDICTIONS

PREDICTION

25–45% CONFIDENCE · BY 2050

By 2050 a skilled person spending a tenth of an average income directs at least 500 HLE a year of accepted AI and robot work.

BASIS

On the central path a year of digital work falls from about \$120 to about \$8.8 by 2050. Income rises about 2.4 times, giving about 660 HLE. The low edge gives only about 120. Early sign: the price per accepted HLE-year of digital work falls below \$35 by 2035.

TRIPWIRE

Proved wrong if, in 2050, a skilled person on a tenth of an average income directs less than 500 HLE a year. A price above \$40 in 2040 would be a strong warning.

QUESTIONS TO ASK

- 1 Which branch is your plan quietly betting on?
- 2 If one person could direct thousands of HLE, what is your organisation for?
- 3 Which checkpoints will you watch, and who owns each one?

15

PART FIFTEEN

WHERE THE WEALTH ACCRUES

Capture: who owns the tools that multiply output when output becomes cheap and plentiful.

15

When one person can direct the work of thousands, who ends up with the money?

Value is gathering at two ends, like the weights on a barbell. Chips, computer power and top AI models sit at one end. At the other end sit people who own trust, customers, data and physical things. Workers' share of income is already slipping.

SCOPE

Compute owners**Scarce assets****Workers**

Earlier chapters kept three ideas apart. Output is how much gets made, productivity is output per hour of work, and capture is who keeps the value. A Roman estate, a medieval mill and a Caribbean sugar farm were all very productive. Yet a few owners kept most of the gain. The twentieth century was different, because workers kept most of the gains from electricity, motors and computers. This chapter asks which pattern the age of AI agents and robots is following. An AI agent is software that can carry out a task on its own, step by step. The early evidence is not comforting.

WHERE OWNERS GAIN

At one end sit the owners of compute (computer processing power) and of frontier models (the most capable AI systems). The biggest cloud firms, often called hyperscalers, planned capital spending that topped \$700 billion for 2026, up from about \$400 billion in 2025 ^[70]. Nvidia, which designs most AI chips, became the first company worth \$5 trillion ^[71].

OpenAI's yearly revenue rate was reported at between about \$50 billion and \$70 billion in autumn 2026. Anthropic's rate passed \$65 billion ^[72]. Together that is many times what the two firms earned two years earlier. These firms are the mill lords of the agent age, owning the machine that everyone else must rent.

WHERE PRICES COLLAPSE

At the other end, the price of machine thinking is collapsing. Running a model as good as GPT-3.5 got more than 280 times cheaper in two years ^[41]. The lower chart shows that fall, starting from 100 in November 2022. The upper chart starts its scale at 50%, so the fall in workers' share looks steeper than it is. Across many tasks, the price of a fixed level of AI skill falls between 9 and 900 times a year. For one common science test it fell about 40 times a year ^[40].


In late 2025 open-weight models, which anyone can download and run, handled about a third of the tokens sent through one large AI marketplace ^[73]. Tokens are small chunks of text that a model reads or writes. Many are Chinese and free, and this competition keeps pushing prices down toward the cost of chips and power.

For the individual, the good news is that the multiplier is cheap to rent and gets cheaper every year. One HLE (Human Labour Equivalent) is one adult's full working year, or 2,500 hours of accepted work. A heavy user on a flat \$180-a-month plan can direct about 18 HLE of accepted AI work. That is the work of about 18 full-time people. It could be up to about 45 if the agents are faster and more of their work is kept. Hiring 18 people would cost about \$2.7 million a year at \$150,000 a head.

The capture barbell

SOURCE: ILO 2024, COMPANY FILINGS 2026, STANFORD AI INDEX 2025, 311 INSTITUTE ANALYSIS 2026

LABOUR SHARE, % (AXIS FROM 50)



MODEL PRICE, NOV 2022 = 100

index



WORKERS' SLIPPING SHARE

Between the two ends sit workers. The labour share is the part of a nation's income paid to workers as wages. Across the world it fell 1.6 percentage points between 2004 and 2024, to 52.3% ^[30].

That small drop is worth about \$2.4 trillion a year, and the ILO (the United Nations' labour agency) links part of it to automation.

Theory gives reasons to expect more. Anton Korinek and Joseph Stiglitz showed in 2017 that AI does not make everyone better off unless the gains are shared ^[74]. Philip Trammell and Korinek showed that if AI can do nearly everything a worker does, labour's share could slide toward zero ^[75]. That could happen even as growth speeds up. Daron Acemoglu and Simon Johnson argue that who gains depends on laws, unions and bargaining power, not on the technology alone ^[76].

CHOKE POINTS IN THE SUPPLY CHAIN

A choke point is a narrow spot that everything must pass through, so whoever controls it can charge a toll or close the road. Our analysis finds three.

The first is magnets. China makes about 94% of the world's sintered permanent magnets ^[59]. These are the strong magnets inside robot joints, electric motors and drones. China curbed exports of rare earths, the metals used in them, in 2025, then paused the curbs for a year from November 2025. That pause runs out around November 2026. More widely, the top three refining countries handle about 86% of key critical minerals ^[92].

The second is advanced chips and the packaging that joins them. TSMC in Taiwan makes most of the world's top AI chips. CoWoS is its way of placing chips and memory side by side on one base. As the chapter on what binds showed, this packaging step is the tightest point today, and TSMC plans to double it by late 2028 ^[91]. Whoever holds that step decides who gets the most capable chips first.

The third choke point is newer: the platforms that take payments for AI agents. When an agent buys, books or sells for you, someone runs the checkout and can take a small cut of every task.

WHERE THE RENTS GO

Economists call income from owning something scarce a rent, and when output is plentiful, rents flow to whatever stays scarce. This report finds six scarce assets that pair with delegated output (work done by AI agents or robots that a person directs).

They are **compute and energy** at scale, and **distribution**, meaning the channels that reach buyers. Next come **trust and proof**, the ability to show what is real, and **brand and identity**. Last are **private data** and **physical assets**, such as land, grid links, factories and mining rights. The chapter on the command equation works out how much one person can direct. There, these assets

shape two of its limits: how much people want to buy, and how much energy there is. Owners of these assets collect rents while everyone else competes on falling prices.

The table below shows layers of the economy and who owns each, not the six assets one by one. A few words in it need a gloss. Capex is capital spending, money spent up front to buy equipment. A "run-rate" is yearly revenue at the current pace. "Commons" means free for anyone to use. "Provenance signing" means content is digitally signed to prove who made it. "Energy-bound" means limited by the supply of power.

LAYER	WHO OWNS IT	EVIDENCE, 2025-26	DIRECTION TO 2100
Compute and chips	Big cloud firms, chip makers	Capex topped \$700bn (2026), Nvidia \$5tn	Concentrated, limited by energy
Frontier models	A handful of labs	Well over \$100bn a year combined revenue run-rate (reported)	Prices fall, open models compete
Open-weight models	Commons, Chinese labs	About a third of routed tokens	Push rents toward zero
Distribution and trust	Platforms, brands, individuals	Provenance signing on by default	Rising share of value
Physical assets	Landowners, grids, miners, robot makers	86% of mineral refining in top 3	Scarce, owners gain
Labour	Workers	52.3% share, falling	Depends on policy and ownership

Source: company filings 2026, ILO 2024, IEA 2025, OpenRouter Dec 2025, 311 Institute analysis 2026.

WHO GAINS POWER, WHO DOES NOT

Our analysis tracks power, meaning the ability to set the terms for others, as well as money. It shows firms and AI systems gaining power over the coming decades, while the individual, as a class, does not. That holds even as each person can make far more than before.

The two facts fit together. A person who directs 1,000 HLE of work still rents the chips, models and payment systems. Big incumbent firms (the large companies already in a market) are better placed to absorb agents than new one-person giants. They already own the customers, the data and the trust. So in the next decade, big firms getting far more output per worker is likelier than millions of solo firms replacing them.

THE INDIVIDUAL'S SHARE

For individuals, the picture cuts both ways. By 2100 a skilled person who spends a tenth of an average income could direct about 262 HLE a year in the setback branch. The figure is about 705 in long drift, 30,000 in agentic diffusion, or 2 million in the explosive window. These are four possible branches, with no odds attached.

But income depends on capture, not on command. A person who directs output that a million others can direct just as easily keeps little of its value. A person who owns something scarce, such as an audience, a trusted name, a data set or a licence, can keep a lot. Income comes to depend on what you own.

THREE WAYS THE GAINS COULD GO

Our analysis gives no single forecast for the labour share, but three capture outcomes. In the first, **owners keep most of the gain**: owners of compute and platforms, and of other scarce assets, take the lion's share. Labour's share could fall below 35% by 2060 in this outcome. Gaps between rich and poor widen sharply, as in the first decades of steam power.

In the second, **gains spread through wages and prices**. Open models and cheap robots push the price of most multipliers toward zero. Rents shrink, buyers gain through cheaper goods, and wages hold up in work that needs a human, such as care, teaching and the trades. This was the twentieth-century pattern.

In the third, **citizens own a share**. Governments spread ownership on purpose, through citizen funds and public compute. A robot or compute dividend is a regular payment to every citizen from robot or compute earnings, not to the owners. Only small examples exist so far. The best known is Alaska, which has paid every resident a share of its oil fund since the 1980s. In 2025 the payment was \$1,000 a person, down from \$1,702 in 2024 ^[122].

These are outcomes, not odds. History shows the first is what happens when nobody acts. The twentieth century shows the second can happen. Alaska shows the third can work, but only on a small scale so far. No G20 country (one of the group of 20 big economies) pays citizens a robot or compute dividend yet.

CHAPTER SUMMARY

Two thousand years of evidence agree. Gains in output go first to whoever owns the multiplier: the slave owner, the mill lord, the factory owner, the platform. In 2025–26 a few chip, cloud and AI firms earn hundreds of billions a year, while machine thinking gets far cheaper. Value moves to what stays scarce.

Cheap output does not make everyone rich. It makes the owners of what stays scarce rich.

CHAPTER CLOSE

WHAT NEXT?

Where the wealth accrues

PREDICTIONS

PREDICTION

25% CONFIDENCE • BY 2035

The share of world income paid to workers as wages falls below 50% by 2035.

BASIS

The ILO put it at 52.3% in 2024, down 1.6 points in 20 years. Early sign: an ILO update below 52%.

TRIPWIRE

Proved wrong if the ILO share is 50% or more in 2035.

PREDICTION

12% CONFIDENCE • BY 2035

By 2035 a G20 economy creates a compute or robot fund paying citizens a regular dividend.

BASIS

Alaska's oil fund pays every resident, but no fund yet pays citizens from compute or robots. Early sign: a G20 budget names one.

TRIPWIRE

Proved wrong if no G20 fund pays such a dividend by 2035.

QUESTIONS TO ASK

- 1 What scarce thing do you own that output cannot copy?
- 2 Which end of the barbell is your business on?

16

PART SIXTEEN

WORK, EARNINGS, TAX AND ECONOMIES

What the new way of producing does to jobs, pay, public money and whole countries.

16

If income used to be hours times a price, what happens to jobs, pay and taxes when hours stop mattering?

The first effect is not mass job loss. It is fewer junior hires. Taxes on wages shrink as output shifts to machines and their owners. Countries split by who owns the scarce parts.

SCOPE

Jobs

Earnings

Tax

Nations

Every institution that touches work assumes the old way of producing. Wages are paid for hours, taxes are charged on wages and pensions grow with years of paid work. Schools and training schemes prepare people to sell their hours. Now one person can direct the work of many AI agents and robots, and that link between hours and income starts to break. This chapter follows the new way of producing into jobs, pay, taxes and whole countries. It also traces the knock-on effects that are easy to miss. The first one shows up in who gets hired at the start of a career.

JOBS: THE JUNIOR SQUEEZE

The first clear effect of AI on jobs is not mass unemployment. Erik Brynjolfsson and colleagues at Stanford studied payroll data. Their study is nicknamed the "canaries" paper, after the birds miners used as an early warning. In August 2025 they found that workers aged 22 to 25 in the most AI-exposed jobs had fallen behind. Their employment was 13% below the path of less-exposed jobs. Their August 2026 update puts the gap at about 19% ^[77]. Older workers held steady.

The IMF (International Monetary Fund) estimates that about 40% of jobs worldwide are exposed to AI, and about 60% in rich economies ^[78]. The World Economic Forum expects 170 million jobs created and 92 million lost by 2030 ^[79]. US employers named AI as a reason for 112,713 job cuts through July 2026 ^[80]. That count is an upper bound, since some firms blame AI for cuts they would have made anyway.

THE BROKEN TRAINING LADDER

This is the pattern of one person plus a swarm of agents, showing up in real data. Firms keep their experienced people, who now direct agents, and stop hiring the juniors whose tasks agents can do. The first-order effect is that output per senior worker rises.

The second-order effect is slower, because junior jobs were where people learned. A trainee lawyer checks contracts for years before she can lead a case. If firms stop hiring juniors, fewer people climb that ladder. Direction capacity is how much delegated work one person can specify, check and stand behind. It is learned on the job, and the ladder that teaches it is being pulled up.

The third-order effect arrives in the 2040s. There may then be too few experts able to check AI work and take the blame when it fails. Checking then becomes the gatekeeper for the whole economy. Firms that kept training juniors will own a scarce skill.

EARNINGS: DIRECTORS AND THE DIRECTED

Pay will split along a new line. On one side are people whose income rises with the output they direct, as long as they capture it.

On the other side are people whose own work can be handed to machines. Their hours now compete with AI work that is very cheap. One HLE of accepted AI work costs about \$120 in 2026, and one HLE is 2,500 hours. So an hour of accepted AI work costs about 5 cents, and that price keeps falling. They also compete with a robot-hour that our analysis puts at \$1 to \$9. Between the two sides sit jobs protected by trust, presence, licences or preference, such as care, teaching, the

skilled trades and many professions. Their pay depends on rules, and on whether people keep wanting a human in the room.

INDIA AND THE PHILIPPINES EXPOSED FIRST

Service exporters sell work done by people at desks to clients abroad. Much of that work is answering customers, writing routine code and processing forms. That is exactly the work AI agents take on first.

In the Philippines, the call-centre and IT services sector employed about 1.9 million people in 2025 ^[100]. It earned more than \$40 billion in export revenue. In India, the tech industry earned about \$315 billion in the year to March 2026, including about \$246 billion from exports ^[121]. It employs about 5.95 million people directly. If agents take over routine service work, that ladder could be pulled up for whole nations, not just for young workers.

POWER PRICES AND FACTORIES COMING HOME

Other knock-on effects run through energy and supply chains. Data centres need huge amounts of power, which can push up everyone's bills. The chapter on what binds showed this in the PJM power auction in the eastern US ^[97]. Data centres added about \$9.4 billion to the cost of that one auction, and those costs end up in household bills. Places that host many data centres may see the same pattern.

Magnets drive another shift. Our analysis expects China's 94% share of strong permanent magnets ^[59] to push other countries to build their own supply. That would bring some factories home, a move called reshoring.

TAX: THE SHRINKING BASE

Most rich-country tax systems lean on workers. Income tax, payroll tax (a tax on wages) and social insurance together raise most of the money. If fewer people direct more agents and robots, and more value goes to owners of capital and compute, the tax base shrinks. That happens just as demands on the state rise.

The table below sets out the problem. The "Today" column shows how each part of the tax system works now. The "With AI and robots" column shows what changes as agents and robots spread, and the "Option" column gives choices for governments.

Bill Gates suggested a robot tax in 2017. South Korea trimmed its tax break for automation the same year. But no major economy taxes robots or compute directly. In one large US cash trial, people got

\$1,000 a month for three years. They worked about 1.3 to 1.4 fewer hours a week, and showed no clear health gains ^[83]. Cash does not replace the meaning, routine and status that work gives.

ISSUE	TODAY	WITH AI AND ROBOTS	OPTION
Tax base	Wages for hours	Output directed, capital and compute	Widen to capital, rents, compute.
Payroll and social insurance	Rises with employment	Erodes as juniors are not hired	Fund from consumption or rents.
Automation	Untaxed, often subsidised	Main source of output	Levy on deployment, or remove subsidies.
Income support	Unemployment insurance	Long transitions, skills gaps	Training accounts, tested cash pilots.
Ownership	Pensions, shares	Multipliers concentrated	Citizen funds, public compute.

Source: OpenResearch 2024, ILO 2024, 311 Institute analysis 2026.

WASHINGTON BLOCKS LEVIES

A levy is a charge or tax set by a government. It could fall on compute, on robots, or on inference, which means running a trained AI model to get answers. A levy on AI compute would mostly tax American firms, and the US has shown it will push back hard. Canada had a digital services tax on big tech firms. In June 2025, after the US halted trade talks over it, Canada dropped the tax a day before the first payments were due ^[98]. That is the main reason our odds are low on a compute levy in an OECD country. The OECD (Organisation for Economic Co-operation and Development) is a club of 38 mostly rich countries.

WORKERS PUSH BACK

Institutions are already responding. Dockworkers on the US East and Gulf Coasts went on strike in 2024. They won a 62% pay rise over six years, plus limits on fully automated equipment ^[81]. A strike by video-game actors in SAG-AFTRA, the US actors' union, ended in 2025. It won rules on consent and disclosure for AI digital copies of performers ^[82]. In May 2026 EU institutions reached a provisional deal to delay the AI Act's high-risk deadlines to about December 2027 and August 2028 ^[84]. Each slows the spread of the multipliers, but none reverses the direction.

COUNTRIES: OWNERS AND RENTERS

Nations will split into owners and renters of the scarce parts. The United States leads in top AI models and cloud computing. China leads in robot making, magnets and mineral refining. One

tracker, IDC, a tech market research firm, says Chinese makers shipped about 78% of humanoid robots in the first half of 2026 ^[53]. (A humanoid robot is shaped roughly like a person.) Another tracker puts China's share above 97%. The Gulf states and Singapore lead in how widely people use AI ^[85].

Europe is rich on the old measures but mostly a renter on the new ones. Poorer countries have a shortcut through phones, open models and falling prices, but many lack computer power, a strong grid and capital. Their share of people using AI tools is falling further behind, not catching up ^[85].

CHAPTER SUMMARY

The evidence points to fewer junior hires, not mass job loss. Young workers in the most exposed jobs are about 19% behind their peers. Pay splits between people who direct output and people whose work is directed. Taxes built on wages face a shrinking base. Countries that own compute, energy and robot supply gain most.

A tax system built on hours sold is built for the old world.

CHAPTER CLOSE

WHAT NEXT?

Work, earnings, tax and economies

PREDICTIONS

PREDICTION

20% CONFIDENCE • BY 2032

By 2032 at least one OECD country brings in a direct levy on AI compute, inference or robot use.

BASIS

The US is the main spoiler. Canada dropped its digital tax under US pressure. Early sign: a finance ministry consults on one.

TRIPWIRE

Proved wrong if no OECD country has passed such a levy by 2032.

PREDICTION

45% CONFIDENCE • BY 2030

By 2030 employment of 22–25-year-olds in the most AI-exposed US jobs trails less-exposed jobs by at least 25%.

BASIS

Stanford's gap was about 19% in August 2026, up from 13% a year before. Early sign: over 22% next update.

TRIPWIRE

Proved wrong if the Stanford gap falls below 10% before 2029.

QUESTIONS TO ASK

- 1 Who will train your experts if junior hiring thins?
- 2 How much of your tax bill rests on payroll?

17

PART SEVENTEEN

MOVES FOR EVERY STAKEHOLDER

What governments, business leaders, investors and individuals should do now.

What should each of us do now that one person's output is no longer one person's worth?

Some moves pay off in every branch of the future. Others are bets on one branch and should wait for a clear signal. Each move needs an owner.

SCOPE

Governments

Executives

Investors

Individuals

This chapter turns the report into action. The chapter on the command equation showed that four things limit how much one person can direct: direction skill, budget, energy and demand. A government, a firm, an investor or a person can change each of them. The chapter on one person in 2100 explains why most of this report's calls rest on one shared bet. The bet is that AI agents become reliable enough to trust without line-by-line checking. The moves below are sorted with that bet in mind. **Safe moves** pay off in every branch of the future. **Contingent moves** are bets on one branch, and should wait until a clear signal, or trigger, appears.

HOW TO READ THE MOVE TABLES

Each stakeholder gets a table with three columns. The first column names the move and the second gives its type. "No-regret" means a safe move that pays in every branch, from setback to explosive window. "Contingent" means a riskier move, and the words after it name the trigger to wait for.

The last column names who should own the move. A move with no owner rarely happens. These safe moves are cheap insurance, while contingent moves only pay if one branch wins. The safe moves that hold in every branch are set out in full in the chapter on one person in 2100.

GOVERNMENTS

The central task for governments is to choose, on purpose, which way the gains should go. The broad sharing of the twentieth century was built through schools, unions, taxes and public ownership, and the agent age needs its own version. That means public compute and data, and schools that teach people to direct work. It also means a tax base that does not rest on hours sold.

MOVE	TYPE	OWNER
Measure delegated output (HLE directed) alongside hours.	No-regret	Statistics office
Protect the junior pipeline: apprenticeships in agent direction.	No-regret	Education and labour ministries
Widen the tax base beyond wages.	No-regret	Treasury
Build public compute and data commons.	No-regret	Science and industry ministry
Secure energy, grid and mineral supply for compute and robots.	No-regret	Energy ministry
Citizen compute or robot dividend.	Contingent: labour share below 50%	Treasury and parliament
Levy on AI inference or robot deployment.	Contingent: payroll base falls 10%+	Treasury

In this table, "HLE directed" means the work a person directs, counted in HLE. A data commons is a store of data that anyone may use. The first trigger is the share of national income paid to workers falling below 50%. The second is the tax raised on wages falling by a tenth or more.

Governments should also watch the choke points. Building some home supply of magnets and chips, or a stockpile, is insurance that pays in every branch.

EXECUTIVES

For business leaders, the old question was how many people they needed. The new question is how much output each person can direct. A second question follows: what scarce thing makes that output worth buying? Revenue per employee will become one of the most watched numbers in business. The firms that win will raise their people's direction skill faster than rivals. They will also own customers, trust or data that free production cannot copy.

MOVE	TYPE	OWNER
Track revenue and output per employee, not headcount.	No-regret	CFO
Raise every employee's direction capacity (k).	No-regret	CHRO and CTO
Keep hiring juniors and redesign their work around agents.	No-regret	CEO
Own a scarce complement: distribution, trust, data or assets.	No-regret	CEO and board
Rent robot-hours before buying fleets.	No-regret	COO
Large humanoid fleet purchases.	Contingent: robots beat humans per task	COO
Rebuild core products as agent-native.	Contingent: 1-month agent horizons	CPO

The owners in this table are job titles. The CEO is the chief executive and the CFO is the finance chief. The CHRO is the people chief, who runs hiring and training, and the CTO is the technology chief. The COO runs day-to-day operations and the CPO is the product chief.

A few other terms need a gloss. Direction capacity, or k, is how much delegated work one person can specify, check and stand behind. A scarce complement is something scarce that goes with cheap output and makes it worth buying. An agent-native product is built from the start for AI agents to use. A time horizon is the length of task, in human working time, that an AI agent can finish on its own half the time. So "1-month agent horizons" means waiting until AI agents can finish month-long tasks.

INVESTORS

For investors, the barbell is the map. Value gathers in compute, energy and the physical supply chain at one end. At the other end it gathers with owners of customers, trust, identity and data. The price collapse hits the middle: firms that sell routine work by the hour, or software that agents can rebuild in a weekend. Tiny teams with huge revenue per person are a sign to watch. The lack of any proven, audited one-person billion-dollar company is a reason for caution.

MOVE	TYPE	OWNER
Weight the barbell: compute and energy, and scarce complements.	No-regret	CIO
Screen for revenue per head and delegated-output intensity.	No-regret	Research head
Underweight hourly-billed and easily rebuilt businesses.	No-regret	Portfolio managers
Hedge supply-chain risk in magnets and minerals.	No-regret	Risk
Back tiny-team and one-person companies at scale.	Contingent: verified \$1bn solo firm	Venture partners
Humanoid platform bets.	Contingent: 1m units a year	Venture and growth

Here the CIO is the chief investment officer, who decides where the money goes. To weight the barbell means putting more money at its two ends. Delegated-output intensity is how much delegated work a firm gets per worker. The two triggers are a verified one-person firm worth \$1 billion, and one million humanoid robots shipped in a year.

INDIVIDUALS

For individuals, the most valuable investment is in their own direction skill. That means the ability to say clearly what work is needed, check it quickly and take responsibility for it. In 2026 a heavy user on a flat \$180-a-month plan can direct about 18 HLE of accepted AI work. For heavy users like this, it is the person, not the money, that limits output.

MOVE	TYPE	OWNER
Learn to direct: specify, verify and own agent work.	No-regret	You
Run agents daily on real work, measure accepted output.	No-regret	You
Build something scarce: audience, reputation, speciality.	No-regret	You
Own a share of the multipliers, not only your hours.	No-regret	You

MOVE	TYPE	OWNER
Go solo with an agent team.	Contingent: you direct 50+ HLE of accepted AI work	You

An ordinary skilled person who spends a tenth of an average income, about \$2,600 a year, gets about 56 HLE in all. Four-fifths of that budget, about \$2,100, buys AI work at about \$120 per HLE. That comes to about 17 HLE. The last fifth, about \$500, buys robot work. Robot work costs about \$15,000 per HLE, so that money buys almost nothing yet.

Add 1 HLE of their own work and about 37 from their share of everyday machines, and the total is about 56. For this person the budget limits output first. In the agentic diffusion branch it stays the limit every decade to 2100. The heavy user spends everything on AI work, so they get 18 HLE for less money. The trigger for going solo is directing 50 HLE or more of accepted AI work, not counting machines. That is well above the 18 a heavy user manages today.

The second investment is in something scarce, such as an audience, a reputation or a speciality. That is what turns cheap output into income.

CHECKPOINTS FOR 2030 AND 2035

Branches are not chosen once but year by year, and a few checkpoints will show which branch we are on. By 2030, watch how many agent sessions top users run at once, and how long a task an AI agent can finish alone. Watch whether employment of young workers in exposed jobs keeps falling behind. By 2035, watch whether leased humanoid robots beat human workers on cost in real warehouses. Watch whether any country has started paying citizens a dividend from compute or robots. Each answer tells us more about which branch we are on.

CHAPTER SUMMARY

Safe moves pay off in every branch, from setback to explosive window. Measure delegated output. Build direction skill. Protect the junior pipeline. Own something scarce. Spread ownership of the multipliers. Prepare taxes for a shrinking wage base. Riskier moves, such as compute dividends, robot levies and big humanoid fleets, should wait for the triggers this report sets out.

The winners will be those who raise their direction skill and own something that stays scarce.

CHAPTER CLOSE

WHAT NEXT?

Moves for every stakeholder

PREDICTIONS

PREDICTION

20% CONFIDENCE · BY 2030

By 2030 at least one G20 national curriculum or major professional licence names and tests agent direction (setting, checking and managing AI work).

BASIS

Firms adopt agents fast, but curricula and licences often take five years or more to change. Early sign: an exam board or law body drafts a standard naming it.

TRIPWIRE

Proved wrong if no G20 curriculum or major licence names and tests agent direction by end-2030.

QUESTIONS TO ASK

- 1 Which safe move will you start this quarter, and who owns it?
- 2 Which trigger would change your plans most if it appeared tomorrow?
- 3 How much AI work can you direct today, and how will you raise it?

18

PART EIGHTEEN

THE OUTPUT DATA BOOK

The full century and decade series, the settings behind the future paths, and how to read each table.

What exactly sits behind every number in this report?

Every figure in the report is worked out once, from the series and settings below. Past values are graded by how good their sources are. Every value after 2026 is an estimate of the future built by our analysis.

SCOPE

Century series

Decade series

Settings

This chapter gives the full data behind the report, so that every figure can be checked, challenged and improved. One set of calculations produces all the series used in earlier chapters. Nothing is typed in twice, so a number in one chapter always matches the same number here. Values before 2026 are graded by source in the source list at the back. Values after 2026 are estimates of the future, not measurements. Each table below comes with a short note on how to read it. Start with the four branches, since every future table uses them, then the units. The century series and decade series follow, and the command equation's settings come last.

THE FOUR BRANCHES, IN BRIEF

The future tables use four named paths. **Setback** is a hard shock in the 2040s, such as a war, pandemic or grid failure, that costs a decade before a slow recovery. **Long drift** is slow, steady change, with AI spreading gradually. **Agentic diffusion** is AI agents and robots spreading widely and doing most routine work. **Explosive window** is very fast growth, where AI speeds up its own progress.

Up to 2050 the last three form a band: long drift is low, agentic diffusion central and explosive window high. After 2050 all four are branches with no odds.

BRANCH LETTERS S, D, A, E

Two tables are too wide for full names, so they use one letter per branch. These are **S** setback, **D** long drift, **A** agentic diffusion and **E** explosive window. So "2100 A" means 2100 in the agentic diffusion branch.

UNITS AND CONVERSIONS

One HLE is one adult working year of output, counted in one of two ways. The physical way is about 75 watts of useful work for 2,500 hours, which is 190 kWh or 0.684 GJ. The time way is 2,500 accepted hours of work by a person, an AI agent or a robot.

Value is always shown in 2025 US dollars and never turned into HLE. PPP dollars are dollars adjusted so they buy the same amount in every country. Old figures in 1990 dollars are multiplied by about 2.25 to reach 2025 dollars. Figures in 2021 PPP dollars are multiplied by about 1.17.

ITEM	VALUE
1 HLE, physical	~75 W × 2,500 h = 190 kWh = 0.684 GJ of useful work
1 HLE, time	2,500 accepted hours of work by a person, agent or robot
Value	Always shown in 2025 US dollars, never converted into HLE
1990 GK\$ to 2025 \$	× ~2.25 (indicative)
2021 PPP \$ to 2025 \$	× ~1.17
Household size	4.0 (2000) to 3.5 (2026) to 3.0 (2100)
Population	8.25bn (2026), 9.7bn (2050), 10.2bn (2100)

Source: Smil 2017, Maddison, World Bank 2026, UN WPP 2024, 311 Institute analysis 2026 (baseline).

CENTURY SERIES, 1 AD TO 2000

The century series tracks Roman Italy up to 400 AD, then Western Europe. Each row is one year, and the columns show GDP per head in 1990 and 2025 dollars. GDP is the value of everything a country makes in a year. The next column gives the value a family of four made in a year, in 2025 dollars. The 1800 row uses Maddison's figure for 1820.

For example, the average Roman family made about \$9,900 of value a year. A Western European family made about \$3,800 in 1000, \$6,900 in 1500 and \$10,800 in 1800. The figure was about \$26,000 in 1900 and \$175,000 in 2000.

The last two columns show useful work, counted in HLE. Useful work is a person's own labour plus energy servants, the animals, machines and fuels that work in place of human muscle. Their energy is divided by 0.684 GJ. The family of four did

about 10 HLE of useful work in 1 AD. It was 8 in 1000, 11 in 1500 and 23 in 1800. It reached 102 in 1900 and 294 in 2000.

All of these are estimates, and values before 1500 can be wrong by a factor of two either way. The second century table names the top multiplier in each century, such as slaves, mills or steam, and who held it.

YEAR	GDP PER HEAD, 1990 \$	GDP PER HEAD, 2025 \$	FAMILY VALUE, 2025 \$	ENERGY SERVANTS PER HEAD	FAMILY USEFUL WORK (HLE)
1	1,100	2,475	9,900	1.9	10
100	1,100	2,475	9,900	1.9	10
200	1,000	2,250	9,000	1.8	10
300	900	2,025	8,100	1.7	9
400	800	1,800	7,200	1.6	9
500	550	1,238	4,950	1.3	8
600	450	1,012	4,050	1.2	7
700	430	968	3,870	1.2	7
800	450	1,012	4,050	1.3	8
900	440	990	3,960	1.3	8
1000	427	961	3,843	1.4	8
1100	520	1,170	4,680	1.5	8
1200	600	1,350	5,400	1.7	9
1300	750	1,688	6,750	1.9	10
1400	790	1,778	7,110	2.1	11
1500	771	1,735	6,939	2.2	11
1600	894	2,012	8,046	2.6	13
1700	997	2,243	8,973	3.0	15
1800	1,202	2,704	10,818	5.0	23
1900	2,890	6,502	26,010	25.0	102
2000	19,456	43,776	175,104	73.0	294

Source: Maddison, Lo Cascio & Malanima, Broadberry, Kander, Malanima & Warde, Morris, 311 Institute analysis 2026 (historical baseline, our estimate).

YEAR	TOP MULTIPLIER	WHO HELD IT	WORLD / FRONTIER NOTE
1	Slaves, oxen, roads, shipping, first watermills	Senatorial and equestrian landowners. The state	Han China ~450. India ~450
100	Imperial trade at its Pax Romana peak	Elite estates, army, state grain supply	Han China ~450
200	Watermills spread (Barbegal)	Estates and the state	Antonine Plague 165–180
300	Tied tenants (colonate)	Great landowners. The taxing state	Price Edict of 301
400	Villa self-sufficiency	Estate lords	Western collapse begins
500	Monastic watermills	Kings, monasteries and the Church	Byzantium higher
600	Watermills under monasteries	Monasteries. Warrior aristocracy	Tang China ~450–600
700	Watermills. Slow recovery	Monasteries. Warrior aristocracy	Tang China the frontier

YEAR	TOP MULTIPLIER	WHO HELD IT	WORLD / FRONTIER NOTE
800	Heavy plough, three-field rotation	Manorial lords. The Church	Abbasid Iraq ~600–700
900	Horse collar spreads	Manorial lords	Abbasid world ~600+
1000	Watermills: 1 per ~200–300 people by 1086	Lords and abbeys (milling monopolies)	Song China ~1,250–1,500, the frontier
1100	Mills, fairs, credit	Lords. Merchant guilds	Song coal and iron
1200	Windmills (from the 1180s)	Lords. Merchant guilds	Song China still ahead
1300	Water-powered fulling and sawing. Compass	Merchant republics. Lords	Italy ~1,400–1,650
1400	Labour scarcity after the Black Death	Surviving peasants gain	Serfdom ends in the West
1500	Printing press. Caravel	Printers, merchants, Crown monopolies	China ~1,000–1,100
1600	Dutch windmills. Fluyt. Joint-stock companies	VOC/EIC shareholders. Planters	Atlantic slavery scales up
1700	Coal. Newcomen engine (1712)	Coal owners. Planters. Slave traders	China ~850. India ~620
1800	Watt engine. Factories. Coke iron	Mill owners. Industrial capitalists	World ~670 (1820). England ~10 energy servants per head
1900	Steam, rail, steel, electricity, chemicals	Industrial firms. Financiers. Empires	World ~1,260. UK ~4,490
2000	Electricity, engines, computers, the internet	Firms, shareholders and, through wages, labour	World ~6,000 (1990 GK\$)

Source: as above (historical baseline).

DECADE SERIES, 2000 TO 2100

The decade series follows the average person on Earth, with measured values to 2026 and estimates from 2030. Realised output adds five parts: the person's own work, physical machine work, robots, digital work and virtual output. Digital work covers software, data, services and content on ordinary screens, including AI-made music, video and text. Virtual output is work made in or for VR, AR and spatial computing, such as immersive worlds and 3D digital twins. It is about 1% of AI work today. The family view multiplies by the average household size, which falls from 4.0 in 2000 to about 3.0 by 2100.

So the family in these tables is the average family on Earth, which made about 125 HLE in 2000. It is not the Western European family of four from the history chapters, which made about 294 HLE in 2000. The two numbers differ because Europe used far more energy per person than the world average.

The first decade table shows the agentic diffusion branch only. Each row is a year. Reading across, the columns give income per head, then the five parts of output, then the total per person and the family figure. It is the average person's total, including their share of everyday machines. In 2100 in this branch, about 76% of it is digital, 19% virtual and 5% physical and own work. The 74% digital, 25% virtual, about 1% physical mix in earlier chapters is different. It is what one skilled person directs.

The second decade table uses the S, D, A, E labels. The first four number columns show income per head in 2025 dollars, and the last four show output per person in HLE. Up to 2050, the D, A and E columns are a low-to-high band. After 2050, all four columns are branches with no odds. Setback only splits away in the 2040s, so before then it tracks long drift closely.

In the agentic diffusion branch, measured income grows more slowly than output. By 2060 output volume has grown about 1.7 times more than measured value, about 6 times against about 4 times since 2026. By 2100 output per person is about 100 times its 2026 level, and income per head about 14 times, a gap of about 7 times. GDP misses much of the cheap new output.

YEAR	GDP/HEAD	OWN	PHYSICAL	ROBOTS	DIGITAL	VIRTUAL	TOTAL/HEAD	AVERAGE GLOBAL FAMILY
2000	15k	0.48	31	0	<0.01	0	31	125
2010	19k	0.47	34	0	<0.01	<0.01	35	133
2020	22k	0.46	35	0	<0.01	<0.01	36	128
2026	26k	0.45	37	<0.01	0.14	<0.01	38	132
2030	29k	0.44	44	<0.01	0.38	<0.01	44	151
2040	41k	0.43	61	<0.01	4.1	0.17	65	216
2050	64k	0.41	79	0.07	43	3.2	126	403
2060	99k	0.39	99	0.37	128	14	241	747
2070	146k	0.37	118	0.94	375	54	548	1.7k
2080	206k	0.36	140	1.7	739	130	1.0k	3.0k
2090	277k	0.34	163	2.8	1.4k	306	1.9k	5.7k
2100	373k	0.32	187	4.2	2.8k	701	3.7k	11k

Agentic diffusion branch, HLE per average person unless stated. Source: 311 Institute analysis 2026 (observed to 2026, our estimates after).

YEAR	INCOME S	D	A	E	OUTPUT S	D	A	E
2000	15k	15k	15k	15k	31	31	31	31
2010	19k	19k	19k	19k	35	35	35	35
2020	22k	22k	22k	22k	36	36	36	36
2026	26k	26k	26k	26k	38	38	38	38
2030	28k	28k	29k	30k	39	39	44	53
2040	32k	32k	41k	53k	41	44	65	118
2050	26k	38k	64k	165k	38	50	126	1.1k
2060	29k	44k	99k	427k	41	56	241	3.3k
2070	34k	51k	146k	841k	44	63	548	11k
2080	39k	59k	206k	1.37m	49	73	1.0k	23k
2090	46k	68k	277k	2.03m	54	86	1.9k	49k
2100	53k	79k	373k	3.00m	62	105	3.7k	106k

S = setback, D = long drift, A = agentic diffusion, E = explosive window. Income is GDP per head in 2025 \$ PPP. Output is HLE per person. Source: 311 Institute analysis 2026 (our estimate).

COMMAND EQUATION SETTINGS

The command series asks how much an ordinary skilled person can direct if they spend a tenth of an average income on delegated work. A fifth of the budget goes on physical work. The virtual share goes on VR, AR and spatial work, and the rest goes on digital work. In the agentic diffusion branch the virtual share is about 1% in 2026, 4% in 2040 and 20% by 2100. The person also directs their share of everyday machines.

In the settings table each row is one setting. Direction capacity, or k, is how much delegated work one person can specify, check and stand behind, in HLE a year. The next two rows give the price of one HLE of digital and robot work, and the last is

the budget share. The columns use the S, D, A, E labels with a year. Note that k caps only delegated work, not the person's share of everyday machines. That is why 56 HLE in 2026 can sit above a k of 20.

PARAMETER	2026	2040 A	2070 A	2100 S	2100 D	2100 A	2100 E
Direction capacity k (HLE)	20	230	14,000	1,000	3,000	110,000	2,000,000
Digital price (\$ per HLE)	120	15	3	20	10	1	0.10
Robot price (\$ per HLE)	15,000	3,000	900	3,000	1,500	500	100
Budget share of income	10%	10%	10%	10%	10%	10%	10%

Source: 311 Institute analysis 2026 (our estimate, calibrated to the 2026 worked case).

In the command table each row is a year. The four branch columns show what one skilled person can direct, in HLE a year. In 2026 it is about 56 in every branch. By 2050 it is 120 at the low end of the band, 661 in the central case and 6,136 at the high end. Setback gives about 59.

The "Limit" column names the limit that binds first in the agentic diffusion branch. It reads "budget" in every row, because the ordinary person runs out of money before running out of time to check work. That holds every decade to 2100. Direction capacity binds first only in the explosive window, and today for heavy users. The last column shows a frontier person with the capital to direct far more.

YEAR	SETBACK	LONG DRIFT	AGENTIC DIFFUSION	EXPLOSIVE WINDOW	LIMIT (DIFFUSION)	FRONTIER (DIFFUSION)
2026	56	56	56	56	budget	58
2030	60	62	80	105	budget	85
2040	73	87	281	694	budget	292
2050	59	120	661	6.1k	budget	1.3k
2060	72	173	1.6k	48k	budget	7.1k
2070	94	259	4.0k	144k	budget	14k
2080	128	358	8.1k	347k	budget	28k
2090	180	499	16k	833k	budget	55k
2100	262	705	30k	2.0m	budget	110k

What one ordinary skilled person can direct, HLE a year. Source: 311 Institute analysis 2026 (our estimate).

CHAPTER SUMMARY

The data book sets out the century series from 1 AD to 2000, the decade series to 2100, the command equation's settings and the units. Past values rest on Maddison, Broadberry, Malanima and Morris, and carry wide error bars. Future values are estimates to be improved.

Every number here can be argued with, and that is why we publish them.

CHAPTER CLOSE

WHAT NEXT?

The output data book

QUESTIONS TO ASK

- 1 Which setting would you change first, and what evidence would justify it?
- 2 Which past value most needs checking before you quote it?
- 3 What data does your organisation hold that could test these estimates?

STRATEGIC SUMMARY

For most of history one body made one body's worth of work. Every time that rule broke, it changed who got rich. Slaves and oxen broke it for Roman landowners. Mills broke it for medieval lords, and coal and steam for mill owners. Fossil-fuelled machines broke it, for once, for nearly everyone in the rich world. AI agents and robots break it again. This time anyone with a phone and a small budget can rent the multiplier. A skilled person in 2026 already directs a small team's work for the price of a subscription. In 2100 a skilled person could direct about 260 HLE a year after a setback, 700 in a long drift, 30,000 if agents spread widely or 2 million in an explosive window. We give no odds that far out. The agentic diffusion branch matches what Isidorus directed with 4,116 slaves and 3,600 pairs of oxen, but nobody is owned.

Three conclusions follow. First, the limit on one person's output moves from the body to the mind. Direction capacity is the skill of setting, checking and standing behind work done by others. It becomes the most valuable skill in the economy. For heavy users today it is already the person, not the computers or the power supply, that limits output. For most people the budget binds first, so the price of AI work matters too. Second, work becomes plentiful and its price falls, so value moves to what stays scarce. That means computers and energy at scale, customers, trust, identity, data and physical assets. Third, who keeps the gains is a choice, not a law of technology. The twentieth century shows that gains can be shared widely. The two thousand years before it show what happens if nobody chooses.

For governments, firms, investors and people, the moves that pay are the same in every branch. Measure delegated work. Raise direction capacity. Protect junior jobs. Own something scarce. Spread ownership of the multipliers. The predictions in this report are not locked. Each one says what early sign to watch and what would prove it wrong, so they can be refined as evidence arrives.

I QUESTIONS TO TAKE AWAY

- ① How much work could you direct today, and what stops you directing more?
- ② Who owns the multipliers in your work: the agents, the robots, the data and the customers?
- ③ Which of your plans still work in all four branches, including the setback?
- ④ What would you measure if hours worked stopped telling you how much a person makes?
- ⑤ How will the next generation learn a craft if junior jobs disappear?

Keep the conversation going. These questions are a starting point for a board, a classroom or a team meeting. The 311 Institute can help turn them into a plan for your organisation.

ASSUMPTIONS & CAVEATS

What to bear in mind

The judgements that most affect how this report should be read, highest impact first. The full set — significant assumptions, market-boundary decisions, source limitations and areas of uncertainty — is set out in the Annex under Assumptions and Caveats.

HIGH IMPACT

Numbers after 2026 are estimates of the future, not measurements.

They rest on assumed paths for prices, direction capacity, energy and robots, and the 2100 branches differ by about 7,600 times.

HIGH IMPACT

Direction capacity (k) matters most and is measured least.

No one measures it yet. The 2026 value is set from one worked case and from what heavy users report.

MEDIUM

Old and very recent figures have wide error bars.

Estimates before 1500 differ by about two times, and several 2025–26 business figures are reported, not audited.

311

SECTION

ANNEX

NOTES AND METHOD

The full method is in Part II, How we built the numbers. This note is a short summary.

Unit: one HLE (Human Labour Equivalent) is one adult working year, counted as about 190 kWh (0.684 GJ) of useful work or as 2,500 accepted hours of work by a person, agent or robot. Money is shown separately in 2025 US dollars and never turned into HLE.

Baseline: a Roman family of four in 1 AD, followed through Europe in 100-year steps to 2000, then an average global family in 10-year steps to 2100.

Three measures are kept apart: output (how much work is done), productivity (output per hour) and capture (who keeps the value).

The Command Equation says output O equals own time h times one plus the work directed in every channel, divided by T (2,500 hours). The work in a channel is units n times hours t times speed s times the share accepted r . It is capped by the smallest of four limits: direction capacity k , budget B divided by price p , energy E divided by energy per HLE e , and demand D .

Up to 2050 each prediction has a sized chance, an early sign and a test that would prove it wrong. After 2050 the report gives four branches with no odds.

Every source was checked again against its origin in October 2026. Two historical figures that could not be checked again are marked in the source ledger.

This is analysis for decision-makers, not investment or legal advice.

ASSUMPTIONS AND CAVEATS

This section states, in full, the assumptions the findings rest on, where the market lines were drawn, the limits of the evidence base, and where the findings are genuinely unresolved — so the reader can assess the conclusions independently. Any sub-section with nothing material to declare says so.

AREAS OF UNCERTAINTY

- In 2100 one skilled person could direct about 260 to 2 million HLE a year across the four branches, and we give no odds after 2050.
- Growth paths run from steady 20th-century growth to explosive growth if most tasks are automated.
- Who keeps the gains depends on choices by governments and owners that the numbers do not model.

MARKET-BOUNDARY DECISIONS

- Supply side only: the report measures how much work gets done, not spending, welfare or happiness.
- Global averages from 2000, and a European spine from 1 AD to 2000 with world comparisons.
- Forced labour is counted as work directed by its owners and named plainly as other people's multipliers.

SIGNIFICANT ASSUMPTIONS

- 1 HLE is about 0.684 GJ of useful work, or 2,500 accepted hours of work. Money is counted separately in 2025 dollars.
- An ordinary person spends 10% of an average income on work they direct. A fifth goes on physical work, the virtual share (about 1% in 2026, 20% by 2100 if agents spread widely) on VR, AR and spatial work, and the rest on digital work.
- World population follows the UN 2024 middle projection, and household size falls from 3.5 to 3.0 people by 2100.
- Only half of all data-centre computing serves directed work in the output series.

SOURCE LIMITATIONS

- Old currency conversions are rough (1990 dollars to 2025 dollars at about 2.25).
- AI task-length scores come from METR software and research tasks, which are close to the limit of what the test can measure.
- Humanoid robot forecasts come from banks and makers that have a business interest in the sector.
- AI company revenue run-rates are reported figures, not audited accounts, and one solo-company case (Medvi) is self-reported and disputed.

TRUST BANDS

PRIMARY

Original / authoritative source.

CORROBORATED

Confirmed from 2+ independent sources.

SINGLE-SOURCE

One source, uncorroborated — hold lightly.

UNVERIFIED

Authenticity or accuracy in question — treat with caution.

FORESIGHT

311 projection / inference — labelled, not fact.

GLOSSARY

Plain meanings for the terms used in this report. Each term is also explained in the text the first time it appears.

01 Acceptance rate

The share of work by an AI agent or robot that passes review and is kept. Rejected work does not count as output.

02 Agent-native

Built from the start around AI agents doing the work, rather than adding agents to an old way of working.

03 Agentic diffusion branch

The branch in which AI agents spread widely and get cheap, and people learn to direct them. Up to 2050 it is the middle of the band.

04 AI agent

Software that can take a goal and carry out a task on its own, step by step, with little help.

05 ARR

Annual recurring revenue: the yearly value of sales that repeat, such as subscriptions.

06 Banalités

Fees that medieval lords charged peasants for using the lord's mill, oven or wine press, which peasants were forced to use.

07 Band

The range of outcomes up to 2050, from a low edge to a high edge, with probabilities attached.

08 Branch

One of four possible paths after 2050. Branches carry no odds, because nobody can honestly size them.

09 Capacity auction

A sale in which power grids pay plants to be ready to supply electricity in future years. Prices rise when spare power is short.

10 Capex

Capital spending: money a firm spends on long-lasting things like buildings, chips and machines.

11 Capture

Who keeps the money that output earns: workers, owners, platforms or the state.

12 Colonnate

A late Roman system that tied tenant farmers to the land they worked, so they could not leave.

13 Command Equation

This report's formula for how much one person can direct. It adds their own work to what their helpers do, capped by the first limit that binds.

14 Compute

Computer processing power, the raw work that chips do to run AI and other software.

15 Critical minerals

Raw materials like lithium, cobalt and rare earths that modern machines need and that come from few places.

16 Data centre

A large building full of computers that store data and run AI and online services.

17 Delegated output

Work that a person hands to AI agents, robots or machines and then directs and checks.

18 Digital output

Everything digital that one person creates: software, apps, data, analysis, digital services and content on ordinary screens, such as text, images, music, video and games. AI-made songs, videos, images and articles count here.

19 Digital twin

A detailed computer copy of a real thing, such as a factory or a city, used to test changes before making them in the real world.

20 Direction capacity

How much delegated work one person can specify, check and stand behind in a year, measured in HLE. It is written k.

21 Early sign

Something we can watch for soon that would show a prediction is on track.

22 Energy servant

One HLE of physical work done by something other than a human body, such as an animal, a mill or a motor.

23 Energy slave

Buckminster Fuller's 1940 term for the work done for each person by machines and fuel. This report's stricter version is the energy servant.

24 Engels' pause

A period in early industrial Britain, about 1780–1840, when output per worker rose much faster than real wages.

25 European spine

Following Europe's history as the main thread of the story. Before 2000 the report follows Roman Italy and then Western Europe, with world comparisons.

26 Explosive growth

Economic growth of more than about 30% a year, far faster than anything seen so far.

27 Explosive window branch

The branch in which AI can do most tasks for a time and growth runs very fast. Up to 2050 it is the high edge of the band.

28 Frontier model

One of the most capable AI models available at a given time.

29 Fully loaded cost

The full cost of employing someone: salary plus tax, pensions, benefits, desk space and other costs.

30 GDP

Gross domestic product: the value of everything a country makes in a year.

31 GDPval

A test from OpenAI that scores AI on real tasks taken from paid jobs.

32 Geary-Khamis dollar

An old unit built to compare buying power across countries and years. Long-run history data use the 1990 version.

33 HLE

Human Labour Equivalent: one adult's working year of output. It is about 75 watts for 2,500 hours, or 2,500 accepted hours of work.

34 Humanoid robot

A robot shaped roughly like a person, so it can use tools and spaces built for people.

35 Inference

Running a trained AI model to get an answer, as opposed to training it.

36 Labour share

The part of a nation's income that is paid to workers as wages, rather than to owners.

37 Lights-out factory

A factory that runs with no workers on the floor, so it can work in the dark.

38 LLM

Large language model: an AI model trained on huge amounts of text to read and write language.

39 Long drift branch

The branch in which AI and robots spread slowly and growth stays modest. Up to 2050 it is the low edge of the band.

40 Malthusian trap

The old pattern where any rise in output led to more people rather than richer people, so living standards stayed flat.

41 METR

Model Evaluation and Threat Research, an independent lab that tests AI systems.

42 MRR

Monthly recurring revenue: the monthly value of sales that repeat, such as subscriptions.

43 Open-weight model

An AI model whose trained settings are published, so anyone can run or change it.

44 Output

How much work gets done, counted in HLE.

45 Pareto-improving

A change that leaves at least one person better off and nobody worse off.

46 Percentage point

The plain difference between two percentages. A rise from 10% to 12% is 2 percentage points (pp).

47 Physical output

Every physical product or physical service, from medicines and drugs to rockets, furniture, food, buildings and moved goods.

48 PPP dollars

Dollars adjusted so they buy the same amount in every country. PPP stands for purchasing power parity.

49 Productivity

Output for each hour worked.

50 Rare-earth magnet

A very strong permanent magnet made with rare-earth metals. Robots, electric cars and wind turbines depend on them.

51 RCT

Randomised controlled trial: a test where people are split by chance into groups, so the effect of one change can be measured fairly.

52 Repo

Short for repository: a store of code for one software project, with its full history.

53 Robot density

The number of industrial robots for every 10,000 factory workers in a country.

54 Run-rate

A firm's current monthly sales multiplied by twelve. It is not audited and can overstate yearly revenue.

55 SaaS

Software as a service: software you rent online by subscription instead of buying and installing it.

56 Serfdom

A medieval system in which peasants were bound to a lord's land and owed him work or payments.

57 Setback branch

The branch in which a hard shock in the 2040s, such as a war, pandemic or grid failure, costs about a decade, then slow recovery.

58 Smart glasses

Glasses with built-in cameras, speakers or small displays, often linked to an AI assistant.

59 Spatial computing

Computing you move around in, using headsets or smart glasses, where digital objects sit in 3D space around you rather than on a flat screen.

60 Teleoperated

Run by a person at a distance, through a remote control link, rather than working alone.

61 Time horizon

The length of a task, measured in human working hours, that an AI agent can finish on its own half the time.

62 Token

A small chunk of text, about three-quarters of a word, that an AI model reads or writes.

63 Tripwire

The test that would prove a prediction wrong. It starts with the words "Proved wrong if".

64 Trust tiers

Grades for sources, in five marks. Primary is direct from the body that measured it. Corroborated means good sources agree. Single-source means only one has it. Unverified means the figure is disputed or self-reported and has not been checked. Foresight is our own estimate about the future. Tables mark the institute's own calculations as "our estimate".

65 Virtual output

Anything made in or for virtual reality, augmented reality, mixed reality and spatial computing. Examples are VR and AR apps, immersive worlds and spatial digital twins.

ABBREVIATIONS

Short forms used in the text, the tables and the list of sources.

01 AI

Artificial intelligence

02 AR

Augmented reality

03 ARR

Annual recurring revenue

04 BCG

Boston Consulting Group

05 BLS

Bureau of Labor Statistics (United States)

06 BofA

Bank of America

07 BPM

Business process management (outsourced office and call-centre work)

08 C2PA

Coalition for Content Provenance and Authenticity

09 CAPEX

Capital expenditure: money spent up front to buy equipment, buildings and chips

10 CEO

Chief executive officer

11 CFO

Chief financial officer

12 CHRO

Chief human resources officer

13 CIO

Chief investment officer (at an investment firm)

14 COO

Chief operating officer

15 CoWoS

Chip-on-Wafer-on-Substrate, an advanced chip packaging method

16 CPO

Chief product officer

17 CTO

Chief technology officer

18 DESA

Department of Economic and Social Affairs (United Nations)

19 EIB

Economic Information Bulletin (a USDA report series)

20 EIC

English East India Company

21 EJ

Exajoule, a billion billion joules of energy

22 ERS

Economic Research Service (part of the USDA)

23 EU

European Union

24 FDA

Food and Drug Administration (United States)

25 FT

Financial Times

26 FY

Financial year

27 G20

Group of Twenty major economies

28 G7

Group of Seven large rich economies

29 GDP

Gross domestic product

30 GJ

Gigajoule, a billion joules of energy

31 GK\$

Geary-Khamis dollar (1990 international dollar)

32 GLP-1

A class of weight-loss and diabetes drugs (glucagon-like peptide-1)

33 H1

First half of the year

34 HBS

Harvard Business School

35 HLE

Human Labour Equivalent

36 IBPAP

IT and Business Process Association of the Philippines

37 IDC

International Data Corporation, a technology market tracker

38 IEA

International Energy Agency

39 IFR

International Federation of Robotics

40 ILA

International Longshoremen's Association, a US dockworkers' union

41 ILO

International Labour Organization

42 IMF

International Monetary Fund

43 ISO

International Organization for Standardization

44 IT-BPM

Information technology and business process management (outsourced tech, office and call-centre work)

45 ITU

International Telecommunication Union

46 kW

Kilowatt, a thousand watts

47 kWh

Kilowatt-hour, a unit of energy

48 LLM

Large language model

49 METR

Model Evaluation and Threat Research, an independent AI testing lab

50 MRR

Monthly recurring revenue

51 MW

Megawatt, a million watts

52 NASSCOM

National Association of Software and Service Companies (India)

53 NBER

National Bureau of Economic Research (United States)

54 OECD

Organisation for Economic Co-operation and Development

55 PJM

The grid operator for much of the eastern United States (PJM Interconnection)

56 pp

Percentage points, written in lower case after a number, as in 19pp

57 PPP

Purchasing power parity

58 Q1

First quarter of the year

59 RCT

Randomised controlled trial

60 SaaS

Software as a service

61 SAG-AFTRA

Screen Actors Guild and American Federation of Television and Radio Artists, a US actors' union

62 TFP

Total factor productivity

63 TSMC

Taiwan Semiconductor Manufacturing Company, the largest contract chipmaker

64 TW

Terawatt, a million million watts

65 TWh

Terawatt-hour, a billion kilowatt-hours

66 UBS

UBS, a large Swiss bank

67 UK

United Kingdom

68 UN WPP

United Nations World Population Prospects

69 USDA

United States Department of Agriculture

70 VOC

Dutch East India Company (Vereenigde Oostindische Compagnie)

71 VR

Virtual reality

72 W

Watt, a unit of power

73 WDI

World Development Indicators (World Bank)

74 WEF

World Economic Forum

75 WEO

World Energy Outlook, the IEA's yearly report

76 XR

Extended reality, a catch-all for VR, AR and mixed reality

SOURCE LEDGER

The report's headline figures, each traced to its origin, dated and tiered, with the full per-claim ledger shipping in the companion data.

#	CLAIM & SOURCE	DATE	TIER
01	Gross domestic product per head by century, 1990 Geary-Khamis dollars, Western Europe 427 (1000), 771 (1500), 1,202 (1820), 2,890 (1900). MADDISON, CONTOURS OF THE WORLD ECONOMY / HISTORICAL STATISTICS (HISTORICAL BASELINE), BOLT & VAN ZANDEN, MADDISON PROJECT DATABASE 2023	2024-04	CORROBORATED
02	Roman Italy GDP per head ~809 (Maddison) to ~1,300–1,800 (Lo Cascio & Malanima) in 1990 dollars. LO CASCIO & MALANIMA, RIVISTA DI STORIA ECONOMICA 25(3) (HISTORICAL BASELINE)	2009-12	SINGLE-SOURCE
03	Roman income: top ~1.5% took ~20% of income, ~90% at or near subsistence, output ~620–680 kg wheat-equivalent per head. SCHEIDEL & FRIESEN, JOURNAL OF ROMAN STUDIES 99 (HISTORICAL BASELINE)	2009-11	SINGLE-SOURCE
04	Slaves in Roman Italy ~1–1.5m of ~6–7.5m people (~15–25%), ~10% empire-wide. SCHEIDEL, JOURNAL OF ROMAN STUDIES 95 (HISTORICAL BASELINE)	2005-11	SINGLE-SOURCE
05	C. Caecilius Isidorus left 4,116 slaves, 3,600 pairs of oxen and 257,000 other animals. PLINY THE ELDER, NATURAL HISTORY 33.135 (PRIMARY ANCIENT SOURCE)	77 AD	PRIMARY
06	Pedanius Secundus kept an urban household of about 400 slaves, all of whom were executed after his murder in 61 AD. TACITUS, ANNALS 14.42–45 (PRIMARY ANCIENT SOURCE)	115 AD	PRIMARY

07	Diocletian's Price Edict: carpenter 50 denarii a day plus food, farm labourer 25. EDICTUM DE PRETIIS RERUM VENALIUM, 301 AD (PRIMARY ANCIENT SOURCE)	301 AD	PRIMARY
08	Sustained human mechanical output ~50–100 W, ox ~300–400 W, horse ~500–750 W. SMIL, ENERGY AND CIVILIZATION, MIT PRESS 2017	2017-05	CORROBORATED
09	Energy capture per head: foragers ~4,000–5,000 kcal/day, Roman West ~31,000 kcal/day, not matched again until ~1700. MORRIS, THE MEASURE OF CIVILIZATION (BASELINE REFERENCE)	2013-01	SINGLE-SOURCE
10	Energy per head: England & Wales ~55–65 GJ c.1800 vs Italy ~18–20 GJ (could not be re-checked in Oct 2026). KANDER, MALANIMA & WARDE, POWER TO THE PEOPLE, MALANIMA, ECHR 69(1) (HISTORICAL BASELINE)	2016-02	SINGLE-SOURCE
11	Domesday Book counts ~5,624 watermills in England, about one per 46 peasant households or roughly one per 200–300 people. DOMESDAY BOOK, 1086 (PRIMARY RECORD), HODGEN, 'DOMESDAY WATER MILLS', ANTIQUITY (1939)	1086	PRIMARY
12	Books printed in Europe in the late 15th century ~12.6m vs ~5m manuscripts for the whole century. BURINGH & VAN ZANDEN, JOURNAL OF ECONOMIC HISTORY 69(2) (HISTORICAL BASELINE)	2009-06	SINGLE-SOURCE
13	Britain ~35,000 hp of steam and ~120,000 hp of water power in 1800, ~2.06m hp of steam by 1870. KANEFSKY & ROBEY (1980), CRAFTS, 'STEAM AS A GENERAL PURPOSE TECHNOLOGY' (HISTORICAL BASELINE)	2004-04	CORROBORATED
14	Atlantic slave trade 1501–1866: ~12.5m embarked, ~10.7m landed. ELTIS & RICHARDSON, TRANS-ATLANTIC SLAVE TRADE DATABASE (SLAVEVOYAGES, LATEST AVAILABLE EDITION)	2023-01	CORROBORATED
15	The US enslaved population ~3.95m in 1860. UNITED STATES CENSUS OF 1860 (PRIMARY RECORD)	1860	PRIMARY

16	Russia's 1861 emancipation freed ~23m privately owned serfs. EMANCIPATION REFORM OF 1861 (HISTORICAL BASELINE)	1861	CORROBORATED
17	Average material living standards in 1800 no better than in the Stone Age (Malthusian frame). CLARK, A FAREWELL TO ALMS, PRINCETON UP (BASELINE REFERENCE)	2007-01	SINGLE-SOURCE
18	Model T chassis assembly fell from 12.5 hours to ~93 minutes on the moving line (1913–14). THE HENRY FORD MUSEUM, MOTORCITIES (HISTORICAL BASELINE)	1914	CORROBORATED
19	The US farm share of the workforce ~41% (1900) to ~1.9% (2000), tractors displaced a working stock of more than 25m horses and mules (1920). DIMITRI, EFFLAND & CONKLIN, USDA ERS EIB-3 (HISTORICAL BASELINE), USDA CENSUS OF AGRICULTURE	2005-06	PRIMARY
20	Containerisation: loading cost ~\$5.86 per ton to ~\$0.16, dock productivity ~1.7 to ~30 tons per labour-hour. LEVINSON, THE BOX, BERNHOFEN, EL-SAHLI & KNELLER (2013) FOR DOCK PRODUCTIVITY (HISTORICAL BASELINE)	2006-01	CORROBORATED
21	Computer performance improved by a factor of between 1.7 trillion and 76 trillion from manual calculation to 2006. NORDHAUS 2007, JOURNAL OF ECONOMIC HISTORY 67(1)	2007-03	PRIMARY
22	Labour hours for 1,000 lumen-hours: ~41.5 (Babylonian lamp), ~5.4 (1800 candle), ~0.0001 (1990s). NORDHAUS, 'DO REAL-OUTPUT AND REAL-WAGE MEASURES CAPTURE REALITY?' (HISTORICAL BASELINE)	1996-01	CORROBORATED
23	In Fortune in 1940 Buckminster Fuller estimated the world had about 36 billion energy slaves, or about 17 per person, with the US holding 54% of them. In 1950 he raised the world average to 38 per person. WIKIPEDIA, ENERGY SLAVE (CITING FULLER, FORTUNE 1940)	1940	SINGLE-SOURCE
24	Output per hour worked 1900–2025: US ~11–12x, UK ~9x, France ~18x (could not be re-checked in Oct 2026). BERGEAUD, CETTE & LECAT, LONG-TERM PRODUCTIVITY DATABASE (LATEST AVAILABLE)	2023-01	SINGLE-SOURCE

25	Annual hours per full-time worker averaged about 3,100 in 1870 (range about 2,750 to 3,480 across 14 countries) versus about 1,350 to 1,800 in Western Europe in 2000. HUBERMAN & MINNS 2007, EXPLORATIONS IN ECONOMIC HISTORY 44(4)	2007-10	PRIMARY
26	World GDP per head about 12.7k (2000), 16.3k (2010), 18.8k (2020) and about 22.5k (2026e), in 2021 PPP dollars. WORLD BANK WDI (NY.GDP.PCAP.PP.KD), IMF WORLD ECONOMIC OUTLOOK APRIL 2026	2026-04	CORROBORATED
27	World primary energy was about 592 EJ in 2024 and passed 600 EJ for the first time in 2025 (up 1.7%), about 73 GJ per head. ENERGY INSTITUTE, STATISTICAL REVIEW OF WORLD ENERGY 2026	2026-07	CORROBORATED
28	World electricity ~30,900 TWh in 2024, ~3,780 kWh per head. EMBER, GLOBAL ELECTRICITY REVIEW 2025	2025-04	CORROBORATED
29	74% of people (~6.0bn) online in 2025, 2.2bn offline. INTERNATIONAL TELECOMMUNICATION UNION (ITU), FACTS AND FIGURES 2025	2025-11	PRIMARY
30	Global labour income share about 52.3% in 2024, down 1.6 points since 2004. At 2024 income levels that fall is worth about \$2.4tn a year of income workers no longer receive. INTERNATIONAL LABOUR ORGANIZATION (ILO), WORLD EMPLOYMENT AND SOCIAL OUTLOOK, SEPTEMBER 2024 UPDATE	2024-09	PRIMARY
31	Global labour share down ~5pp since 1980. KARABARBOUNIS & NEIMAN, QUARTERLY JOURNAL OF ECONOMICS (BASELINE)	2014-02	CORROBORATED
32	Global AI computing capacity doubling about every 7 months. EPOCH AI, DATA INSIGHT 'GLOBAL AI COMPUTING CAPACITY IS DOUBLING EVERY 7 MONTHS'	2026-01	PRIMARY
33	Generative AI assistant: +14% issues resolved per hour, +34% for novices. BRYNJOLFSSON, LI & RAYMOND, QUARTERLY JOURNAL OF ECONOMICS	2025-02	PRIMARY

34	Professional writing with ChatGPT: time -40%, quality +18%. NOY & ZHANG, SCIENCE (BASELINE)	2023-07	PRIMARY
35	Consultants at BCG: +12.2% tasks, 25.1% faster, >40% quality inside the frontier, 19pp less accurate outside it. DELL'ACQUA ET AL., HBS/BCG (BASELINE)	2023-09	PRIMARY
36	Experienced developers 19% slower with early-2025 AI tools while believing they were ~20% faster. MODEL EVALUATION AND THREAT RESEARCH (METR), ARXIV 2507.09089	2025-07	PRIMARY
37	Danish administrative data: no measurable effect of AI chatbots on earnings or hours. HUMLUM & VESTERGAARD, NBER	2025-05	CORROBORATED
38	50% task-completion time horizon: GPT-5 about 2h17m (Aug 2025), Opus 4.5 about 4h49m (Dec 2025), Opus 4.6 about 12h on METR's later rescoring (first reported at 14.5h). The 80% horizon is about ten times shorter, and METR warns its task suite is close to saturated. MODEL EVALUATION AND THREAT RESEARCH (METR), TIME HORIZON 1.1	2026-05	CORROBORATED
39	GDPval: frontier models win or tie against experts on ~48% of tasks (Sep 2025) rising to ~71% (Dec 2025). OPENAI GDPVAL, GPT-5.2 LAUNCH REPORTING	2025-12	CORROBORATED
40	Price of a fixed level of LLM capability falls 9x–900x a year depending on the task (about 40x a year for GPT-4-level PhD science questions). EPOCH AI, LLM INFERENCE PRICE TRENDS	2025-03	PRIMARY
41	GPT-3.5-level inference fell >280x in price, Nov 2022 to Oct 2024. STANFORD HAI, AI INDEX 2025	2025-04	CORROBORATED
42	Base44, solo-founded by Maor Shlomo, sold to Wix for \$80m cash about six months after founding, by which time it had about 8 staff. TECHCRUNCH	2025-06	CORROBORATED

43	Lovable: ~\$400m ARR with ~146 staff (~\$2.7m per employee). TECHCRUNCH, STARTUP PRESS REPORTS	2026-03	CORROBORATED
44	Cursor passed \$2bn annualised revenue in March 2026. BLOOMBERG, VIA TECHCRUNCH	2026-03	CORROBORATED
45	Solo founder portfolio ~\$3m a year (Photo AI ~\$132k MRR). INDIE HACKERS CASE STUDY	2025-12	CORROBORATED
46	Today's agents and robots could in theory automate ~57% of US work hours, ~\$2.9tn a year of value by 2030. MCKINSEY GLOBAL INSTITUTE, 'AGENTS, ROBOTS, AND US'	2025-11	CORROBORATED
47	15% of day-to-day work decisions made autonomously by 2028, >40% of agentic projects cancelled by end-2027. GARTNER	2025-06	CORROBORATED
48	Artificial intelligence adds no more than about 0.66% to TFP over ten years. ACEMOGLU, 'THE SIMPLE MACROECONOMICS OF AI', ECONOMIC POLICY 40(121)	2025-01	PRIMARY
49	Generative AI could lift global GDP ~7% over ten years. GOLDMAN SACHS RESEARCH (BASELINE)	2023-03	PRIMARY
50	Median Gemini text prompt ~0.24 Wh, 33x less energy than a year earlier. GOOGLE CLOUD	2025-08	CORROBORATED
51	Data-centre electricity about 415 TWh in 2024, around 945 TWh by 2030 and about 1,200 TWh by 2035 in the base case, with scenarios spanning roughly 700 to 1,700 TWh in 2035. INTERNATIONAL ENERGY AGENCY (IEA), ENERGY AND AI	2025-04	PRIMARY
52	~542,000 industrial robots installed in 2024 and operational stock ~4.66m. Korea robot density was 1,012 per 10,000 workers in 2023. INTERNATIONAL FEDERATION OF ROBOTICS (IFR), WORLD ROBOTICS 2025 (INSTALLS, STOCK) AND WORLD ROBOTICS 2024 (KOREA DENSITY)	2025-09	CORROBORATED

53	Humanoid shipments were ~13,000 (Omdia) to ~18,000 (IDC) in 2025. IDC puts H1 2026 at ~25,000 with Chinese makers ~78%. Another tracker cited by TechNode puts H1 2026 at ~19,100 with Chinese makers >97%.	2026-09	CORROBORATED
	IDC, OMDIA, TECHNODE CITING IT HOME		
54	1X NEO consumer humanoid at \$20,000 or \$499/month, harder tasks teleoperated.	2025-10	CORROBORATED
	1X / BUSINESSWIRE, PRESS REPORTS		
55	~3bn humanoids by 2060, unit cost ~\$13k–17k by 2030–35.	2026-03	PRIMARY
	BANK OF AMERICA INSTITUTE, PHYSICAL AI PART 2		
56	Humanoid market ~\$5tn by 2050.	2025-05	CORROBORATED
	MORGAN STANLEY RESEARCH		
57	~300m humanoids in use by 2050.	2025-06	CORROBORATED
	SWISS BANK UBS		
58	Humanoid robot market ~\$38bn by 2035.	2024-02	PRIMARY
	GOLDMAN SACHS RESEARCH		
59	China ~94% of sintered permanent magnet output (2024). Export curbs paused for one year from Nov 2025.	2025-11	CORROBORATED
	INTERNATIONAL ENERGY AGENCY, PROJECT BLUE		
60	Standard ISO 25785-1 for dynamically stable (legged) industrial robots at Committee Draft stage.	2025-10	PRIMARY
	INTERNATIONAL ORGANIZATION FOR STANDARDIZATION		
61	Waymo ~500,000 paid rides a week, targeting 1m by end-2026.	2026-03	CORROBORATED
	WAYMO, VIA INSIDEEVS (500K FIGURE), BLOOMBERG (1M TARGET)		
62	Xiaomi's lights-out smart factory (opened Jul 2024) runs 24/7 and can make about one phone a second at full capacity.	2024-07	CORROBORATED
	XIAOMI, NEW ATLAS, BGR		

63	John Deere targets a fully autonomous corn and soybean system by 2030. PRECISION FARMING DEALER	2022-05	CORROBORATED
64	Suno generates a Spotify catalogue's worth of music (~100m tracks) every two weeks. MUSIC PRESS: DJ MAG, EDM.COM	2025-11	CORROBORATED
65	~99,000 new tracks uploaded to streaming services a day in 2024, down from ~103,500 in 2023. LUMINATE 2024 YEAR-END REPORT, MUSIC BUSINESS WORLDWIDE	2025-01	CORROBORATED
66	About half of new web articles are AI-generated, though detection is imperfect and the share has levelled off. GRAPHITE STUDY, REPORTED BY AXIOS (MAY 2026)	2026-05	CORROBORATED
67	Global games market ~\$188.8bn in 2025. NEWZOO	2025-09	CORROBORATED
68	Creator economy ~\$250bn (2023) heading to ~\$480bn by 2027. GOLDMAN SACHS RESEARCH	2023-04	PRIMARY
69	Google Pixel 10 signs camera images with C2PA Content Credentials by default. C2PA / GOOGLE	2025-08	CORROBORATED
70	Big Tech capital spending planned for 2026 topped \$700bn. COMPANY Q1 2026 EARNINGS, YAHOO FINANCE	2026-05	CORROBORATED
71	Nvidia the first company to reach \$5tn market capitalisation. MARKET REPORTS	2025-10	CORROBORATED
72	OpenAI's annualised run-rate is disputed: about \$70bn in Sep 2026 per Reuters, nearer \$50bn per the FT, with about \$70bn expected by year-end (Bloomberg). Anthropic's annualised run-rate topped \$65bn by autumn 2026. Reported, not audited. REUTERS, FINANCIAL TIMES, BLOOMBERG, BLOOMBERG GOVERNMENT	2026-10	UNVERIFIED

73	Open-weight models handle about one-third of tokens routed on OpenRouter, and Chinese open models rose from about 1% to peaks near 30% of weekly tokens in 2025.	2025-12	CORROBORATED
	OPENROUTER WITH A16Z, STATE OF AI: A 100-TRILLION-TOKEN STUDY		
74	Without redistribution AI is not Pareto-improving.	2017-12	PRIMARY
	KORINEK & STIGLITZ, NBER WORKING PAPER W24174 (2017)		
75	If AI substitutes closely for labour the labour share can trend toward zero as growth accelerates.	2025-09	PRIMARY
	TRAMMELL & KORINEK, NBER W31815 (REVISED)		
76	Who gains from technology depends on institutions and bargaining power.	2023-05	PRIMARY
	ACEMOGLU & JOHNSON, POWER AND PROGRESS (BASELINE)		
77	Employment of 22–25 year-olds in AI-exposed occupations is now about 19% below where it would be had it kept pace with less-exposed peers, up from the 13% relative decline first reported in August 2025.	2026-08	CORROBORATED
	BRYNJOLFSSON, CHANDAR & CHEN, STANFORD DIGITAL ECONOMY LAB (AUGUST 2026 UPDATE)		
78	~40% of global jobs exposed to AI, ~60% in advanced economies.	2024-01	CORROBORATED
	INTERNATIONAL MONETARY FUND		
79	170m jobs created and 92m displaced by 2030 (net +78m).	2025-01	PRIMARY
	WORLD ECONOMIC FORUM, FUTURE OF JOBS REPORT 2025		
80	112,713 US job cuts citing AI through July 2026.	2026-08	CORROBORATED
	CHALLENGER, GRAY & CHRISTMAS		
81	Dockworkers of the ILA union won ~62% wage rise over six years plus limits on fully automated equipment.	2025-01	CORROBORATED
	MARITIME EXECUTIVE, BIG THINK		
82	The SAG-AFTRA video-game strike ended with consent and disclosure terms for AI replicas.	2025-07	CORROBORATED
	PRESS REPORTS		

83	\$1,000/month for three years: work ~1.3–1.4 hours/week lower, no clear health gains. OPENRESEARCH / VIVALT ET AL., NBER	2024-07	CORROBORATED
84	EU institutions reached a provisional deal on 6 May 2026 to delay the AI Act's high-risk deadlines to about Dec 2027 (Annex III) and Aug 2028 (Annex I). GIBSON DUNN, ONETRUST	2026-05	CORROBORATED
85	Generative AI used by ~17.8% of the working-age world (Q1 2026), North-South gap widening. MICROSOFT, AI DIFFUSION REPORT 2026 Q1	2026-05	PRIMARY
86	World population 8.2bn (2024), peaking ~10.3bn in the mid-2080s, ~10.2bn in 2100. UNITED NATIONS DESA, WORLD POPULATION PROSPECTS 2024	2024-07	CORROBORATED
87	Explosive growth (more than 30% a year) by 2100 is plausible and cannot be ruled out. DAVIDSON, OPEN PHILANTHROPY (BASELINE)	2021-06	PRIMARY
88	Under default GATE parameters growth passes 30% a year once roughly 50–70% of tasks are automated, well before full automation. EPOCH AI, GATE MODEL (MARCH 2025) AND 'AI AND EXPLOSIVE GROWTH REDUX' (POTLOGEA & HO)	2025-06	PRIMARY
89	Of seven singularity tests only two point toward one, on trend a century or more away. NORDHAUS, NBER W21547 (2015), PUBLISHED IN AEJ: MACROECONOMICS (2021)	2021-01	CORROBORATED
90	At 2.3% a year energy growth, waste heat would boil Earth's surface within ~400 years. MURPHY, DO THE MATH, 'GALACTIC-SCALE ENERGY' (BASELINE)	2011-07	SINGLE-SOURCE
91	Taiwan's TSMC reportedly plans to double CoWoS advanced-packaging capacity from about 130,000 wafers a month at end-2026 to about 260,000 by end-2028. TRENDFORCE, CITING ECONOMIC DAILY NEWS SUPPLY-CHAIN REPORTS	2026-09	CORROBORATED

92

Top-three refiners hold ~86% of key critical-mineral refining (2024).

INTERNATIONAL ENERGY AGENCY, GLOBAL CRITICAL MINERALS OUTLOOK 2025

2025-05

CORROBORATED

93

The HLE, energy-servant and Command Equation constructs and all 2030–2100 values.

311 INSTITUTE ANALYSIS 2026 (CONSTRUCTS, ESTIMATE / FORESIGHT)

2026-10

FORESIGHT

94

Medvi, a GLP-1 telehealth marketer, told the New York Times it made about \$401m in sales in 2025 with essentially two staff (the founder and his brother), using AI tools and outsourcing doctors, prescriptions and pharmacy work to partners such as OpenLoop and CareValidate. The figures are self-reported and unaudited. Critics including Futurism and Telecare Aware said the profile was too flattering, and the Times later added an editor's note about an FDA warning letter and a class action.

NEW YORK TIMES PROFILE, DEALROOM SUMMARY, FUTURISM AND TELECARE AWARE CRITIQUES

2026-04

UNVERIFIED

95

Deezer says fully AI-generated tracks passed 50% of daily uploads for the first time in June 2026 (up from 44% in April 2026), yet AI music still makes up only about 1–3% of streams, and up to 85% of streams on AI tracks were flagged as fraudulent.

DEEZER NEWSROOM

2026-07

PRIMARY

96

Robert Allen calculated that British output per worker rose 46% between 1780 and 1840 while real wages rose only 12%, a gap he called Engels' pause.

R.C. ALLEN, ENGELS' PAUSE, EXPLORATIONS IN ECONOMIC HISTORY

2009-10

PRIMARY

97

PJM's capacity auction price for 2025/26 jumped about ninefold to \$269.92 per MW-day from \$28.92 the year before. PJM blamed mainly plant retirements, market rule changes and higher forecast demand, and its market monitor later estimated that data-centre load added about \$9.4bn to that auction's cost.

PJM RESULTS (POWER MAGAZINE), PJM MARKET MONITOR (BENTON INSTITUTE)

2024-07

CORROBORATED

98

Canada announced on 29 June 2025 that it would rescind its digital services tax, a day before first payments were due, after President Trump suspended trade talks over it.

DEPARTMENT OF FINANCE CANADA AND AL JAZEERA

2025-06

PRIMARY

99

China's State Council AI+ action plan, issued in August 2025, targets more than 70% adoption of new-generation intelligent terminals and AI agents by 2027 and more than 90% by 2030.

STATE COUNCIL OF CHINA VIA SESEC

2025-08

CORROBORATED

100

The Philippine IT-BPM sector earned more than \$40bn in export revenue in 2025, up about 5%, and employed about 1.9 million people, up from 1.82 million in 2024, according to industry body IBPAP.

IBPAP VIA PHILIPPINE DAILY INQUIRER AND CONTEXT.PH

2026-01

CORROBORATED

101

Google said its products processed more than 1.3 quadrillion tokens a month (October 2025), up from about 980 trillion in June 2025. Much of the rise reflects reasoning models that use far more tokens per request, so token counts track computing load more than useful work.

THE DECODER, REPORTING SUNDAR PICHAI AT A GOOGLE CLOUD EVENT

2025-10

CORROBORATED

102

China's GDP per head peaked under the Northern Song at roughly \$1,250–1,500 (1990 international dollars) around 980–1020, above Western Europe at the time.

BROADBERRY, GUAN & LI 2018, JOURNAL OF ECONOMIC HISTORY 78(4)

2018-12

PRIMARY

103

Northern Song iron output estimated at about 125,000 short tons (over 100,000 tonnes) a year by 1078.

HARTWELL 1962, JOURNAL OF ASIAN STUDIES 21(2)

1962-02

SINGLE-SOURCE

104

Barbegal mill complex near Arles had 16 waterwheels. A 2018 study suggests it mostly made ship's biscuit and stood idle for months each year, so older estimates of feeding ~10,000 people are uncertain.

SÜRMELIHINDI ET AL., SCIENCE ADVANCES (VIA POPULAR ARCHAEOLOGY)

2018-09

SINGLE-SOURCE

105

Reconstruction of per-head GDP in central-northern Italy from 1300, high in the Renaissance then falling. Used for the Italian trading-city estimate of about \$1,400–1,650 (1990 \$) around 1300.

2011-08

SINGLE-SOURCE

MALANIMA 2011, EUROPEAN REVIEW OF ECONOMIC HISTORY 15(2)

106

Global XR device shipments including smart glasses grew 44.4% in 2025, driven mainly by smart glasses, with Meta holding 72.2% of the market while its Quest headset shipments fell 42.3%.

2026-04

CORROBORATED

IDC HEADSET TRACKER (ICTBUSINESS)

107

About 2.25 million smart glasses shipped worldwide in Q1 2026, up 167% year on year, with Meta holding 69.2% share.

2026-07

PRIMARY

IDC AR/VR AND WEARABLES TRACKERS

108

Market tracker IDC forecasts about 13.6 million smart glasses in 2026 rising to 27.3 million by 2030, and mixed reality headsets rising from 3.2 million in 2026 to 10.4 million in 2030.

2026-07

PRIMARY

IDC

109

Global AI glasses shipments rose 263% year on year in H1 2026, 96% were display-less models and Meta held 94% of the display-less segment.

2026-09

CORROBORATED

COUNTERPOINT RESEARCH INTELLIGENT EYEWEAR MARKET TRACKER VIA ANDROID AUTHORITY

110

EssilorLuxottica sold more than 7 million Meta smart glasses in 2025, against about 2 million sold from late 2023 to early 2025.

2026-02

CORROBORATED

ESSILORLUXOTTICA Q4 2025 RESULTS VIA ROAD TO VR

111

Meta Reality Labs had 2025 revenue of about \$2.21 billion and an operating loss of \$19.19 billion, up from a \$17.73 billion loss in 2024.

2026-01

CORROBORATED

META Q4 2025 EARNINGS VIA AUGANIX

112

Reality Labs lost \$4.62 billion on \$431 million revenue in Q2 2026, versus a \$4.53 billion loss on \$370 million a year earlier.

2026-07

CORROBORATED

META Q2 2026 EARNINGS VIA GADGETREVIEW

113

Market tracker IDC estimates Apple shipped about 390,000 Vision Pro units in 2024 and only about 45,000 in Q4 2025, and says production was halted at the start of 2025, which is disputed.

IDC ESTIMATE VIA SLASHDOT

2026-01

SINGLE-SOURCE

114

Samsung launched the Galaxy XR headset, the first Android XR device, at \$1,800 in October 2025.

GSMARENA

2025-10

CORROBORATED

115

Carmaker BMW is industrialising its Omniverse-based Virtual Factory with digital twins of more than 30 production sites, projecting up to 30% lower production planning costs.

BMW GROUP VIA NEWMOBILITY.NEWS

2025-06

CORROBORATED

116

Roblox paid creators \$477 million through Developer Exchange in Q4 2025, up 70% year on year, though Roblox is mainly a flat-screen 3D platform not VR.

ROBLOX Q4 2025 RESULTS VIA QUARTR

2026-02

CORROBORATED

117

Roblox creators earned more than \$1 billion from March 2024 to March 2025, up 31%, with the top 10 creators averaging about \$33.9 million.

ROBLOX ANNUAL ECONOMIC IMPACT REPORT VIA OUTLOOK RESPAWN

2025-11

CORROBORATED

118

Google opened Project Genie, built on the Genie 3 world model, to US Google AI Ultra subscribers in January 2026, generating interactive worlds limited to 60 seconds.

GOOGLE BLOG

2026-01

PRIMARY

119

World Labs launched Marble commercially in November 2025. It generates 3D worlds that export as Gaussian splats or meshes and can be viewed in VR on Vision Pro and Quest 3, priced from free to \$95 a month.

RADIANCE FIELDS

2025-11

CORROBORATED

120

Roblox released and open-sourced Cube 3D, a generative AI model that creates 3D objects from text prompts, in March 2025.

ROBLOX NEWSROOM

2025-03

PRIMARY

121

India's tech industry revenue was put at about \$315bn in FY2026, up 6.1%, with about \$246bn of exports and about 5.95 million direct employees, up from about 5.82 million.

NASSCOM STRATEGIC REVIEW 2026 VIA ETV BHARAT

2026-02

CORROBORATED

122

Alaska's Permanent Fund Dividend was \$1,000 per eligible resident in 2025 (paid from 2 October to more than 600,000 people), down from \$1,702 in 2024. Since a 2017 court ruling the amount is set each year by the legislature.

ALASKA'S NEWS SOURCE

2025-09

CORROBORATED

TERMS & CONDITIONS

This report is published by the 311 Institute for information only. It is made available to clients and to the general public alike, and is not advice of any kind.

No advice. Nothing in this report constitutes legal, financial, investment, tax or other professional advice, nor any recommendation, offer or solicitation to buy, sell or hold any security, instrument or asset, and it takes no account of any reader's particular circumstances. Any action taken on the basis of this report is taken entirely at the reader's own risk, and independent professional advice should be obtained before acting.

Foresight, not fact. This report contains analysis, scenarios and forward-looking judgements about uncertain futures. Projections, ranges and scenarios are informed estimates, not statements of fact or guarantees, and actual events may differ materially. Forward-looking statements speak only as at the date of publication and are not updated.

No liability. To the fullest extent permitted by law, the 311 Institute, its founder Matthew Griffin, and their officers, employees, agents and associates accept no liability of any kind — whether in contract, tort (including negligence), for misrepresentation, or otherwise — for any loss or damage whatsoever (including direct, indirect, special, consequential or economic loss) arising from any use of, or reliance on, this report or its contents by any person, whether a client or a member of the public. No representation or warranty, express or implied, is given as to the report's accuracy, completeness, currency or fitness for any purpose. Nothing in these terms excludes any liability that cannot be excluded by law.

Reproduction and intellectual property. All content, frameworks, methods and marks in this report are the property of the 311 Institute and its licensors. No part of this report may be reproduced, copied, stored, distributed, published or otherwise used, in whole or in part, without the prior written permission of the 311 Institute. Third-party names, standards, products and marks are the property of their respective owners and are used for identification only.

Governing law. These terms, and any dispute or claim arising out of or in connection with them or this report, are governed by the law of England and are subject to the exclusive jurisdiction of the English courts. © 2026 311 Institute. All rights reserved.

311 INSTITUTE · INTELLIGENCE UNIT

To help humanity build a future where people and planet thrive, and where the benefits of human progress are shared by everyone — not just the few.



enquiries@311institute.com · 311institute.com · London

© 2026 311 Institute · All rights reserved.