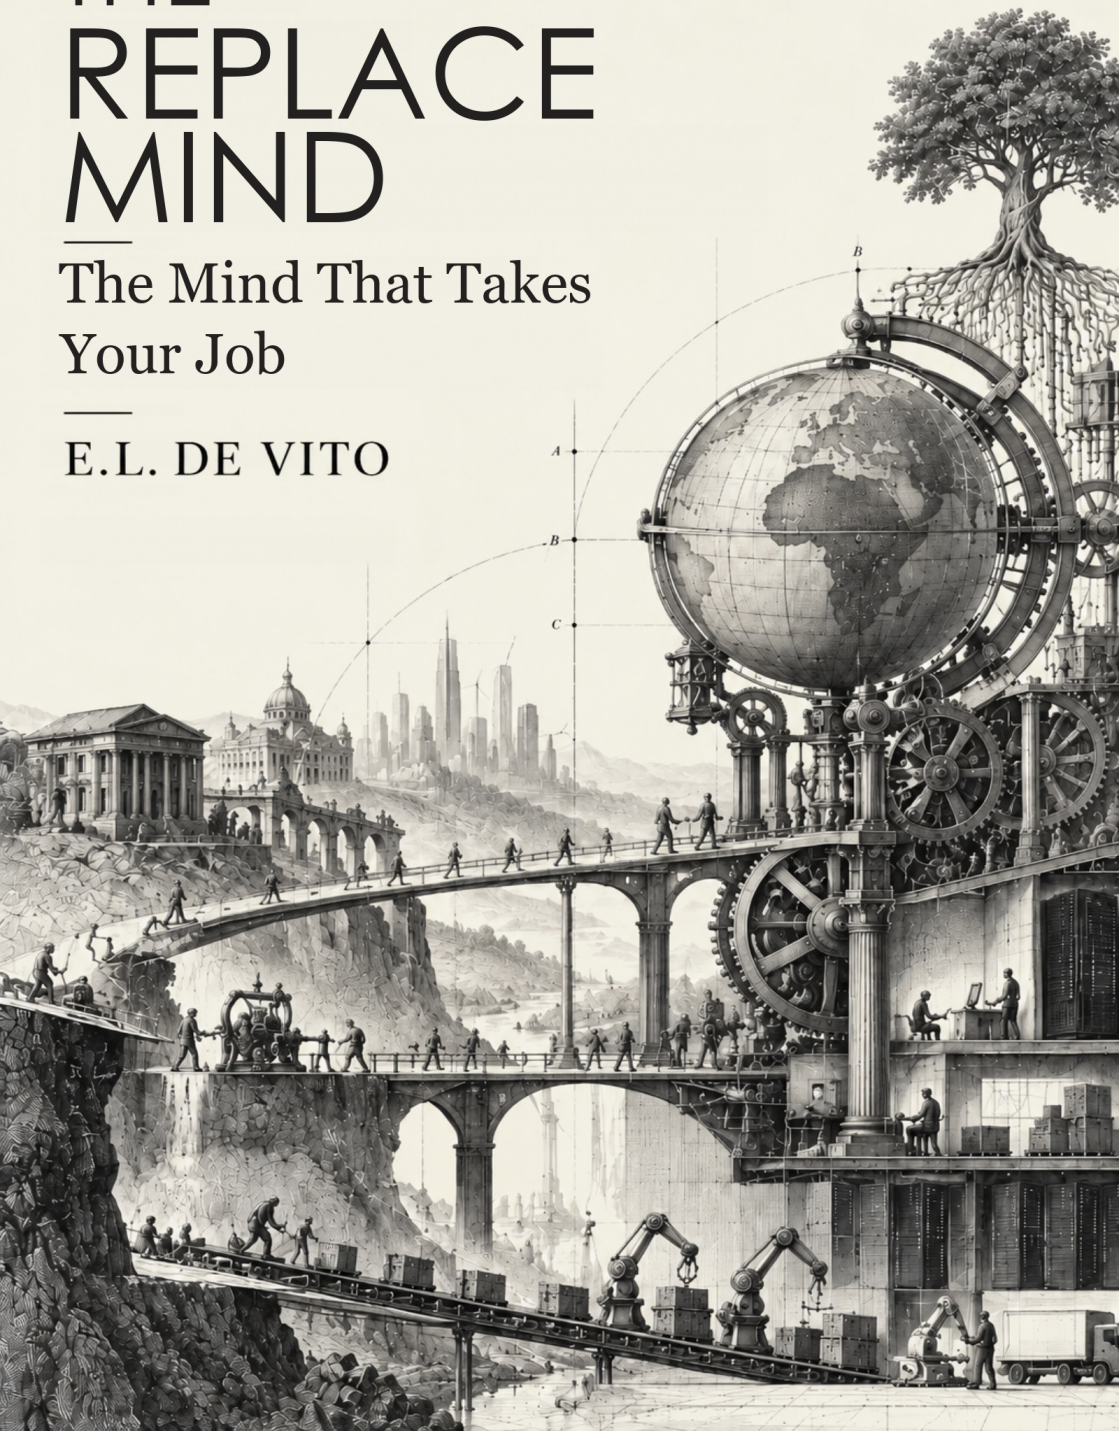
THE REPLACE MIND

The Mind That Takes Your Job

E.L. DE VITO



The ReplaceMind

The Mind That Takes Your Job

by E.L. De Vito

This book was developed with the support of Claude, ChatGPT and Grok; parts of the text and of the code were written with the help of artificial intelligence.

Contents

Introduction....5
Chapter 1 Output, Employment and Income Distribution....9
Chapter 2 Tasks, Exposure and AI Adoption....19
Chapter 3 Participation, Unemployment and Discouragement....29
Chapter 4 Labor Demand and Supply....40
Chapter 5 Labor Income, Capital Income and Consumption....49
Chapter 6 The Model and the Paired Simulations....58
Chapter 7 Outflows and Re-employment....68
Chapter 8 Entering Work for the First Time....76
Chapter 9 Occupational Retraining....86
Chapter 10 Value Added, Income and Regions....97
Chapter 11 Public Revenue and Debt....106
Chapter 12 Calibration, Parameters and Reproducibility....115
Chapter 13 Public Policy Evaluation....132
Chapter 14 Aggregate Results and Heterogeneity....142
Chapter 15 Revising the Model....151
Chapter 16 Sensitivity Analysis....160
Chapter 17 Social Cohesion and Institutions....167
Chapter 18 Model Limits and Empirical Testing....176
Chapter 19 Review of the Evidence....185
Chapter 20 Tax Bases and Public Finance....198
Chapter 21 Scenarios to 2036....204
Chapter 22 Interventions, Responsibilities and Timing....220
Chapter 23 The Causal Structure of the Model....229

Chapter 24 Training and Learning 236
Chapter 25 Job Quality and Pay 246
Chapter 26 AI Adoption in Small Firms 254
Chapter 27 Capital Ownership and Concentration 265
Chapter 28 AI and Health Services 273
Chapter 29 The Energy Use of AI 281
Chapter 30 Mobility and Territorial Differences 290
Chapter 31 Impact Evaluation of Interventions 300
Chapter 32 Results and the Limits of Knowledge 307
Chapter 33 Summary of Results: Work and Incomes 312
Chapter 34 Feedback Between Employment and Incomes 316
Chapter 35 Summary of Policies and Open Choices 319
Chapter 36 The Last Forecast 323
Chapter 37 Social Risk 339
Conclusions Results and Open Questions 349
Chapter Notes 353
References 362
Appendix: The Formulas Explained Step by Step 371
Symbols and Definitions 394
The Code Behind the Book 395

Introduction

They will tell you it has happened before.

That every time a machine took a person's place, someone announced the end of work, and work came back. Mechanical looms, tractors, computers: each time the fear, each time the recovery.

In the long run, that is true. But the people who tell you so usually forget two things.

The first is how long the long run lasted. In England in the first decades of the nineteenth century, output per worker grew year after year while real wages stood still. For thirty years. An entire generation worked inside the Industrial Revolution without seeing its fruits. Work did come back, yes: to their children, in other cities, in other trades. Not always to the people who had lost it.

The second thing matters even more. Each time there was somewhere to go. Those who left the fields found the factory. Those who left the factory found the office, the counter, the service desk. Work done with the head was the refuge: the part that machines could not do.

This time the machine walks straight into that refuge. It reads, writes, classifies, answers, checks. And together with robots it is beginning to approach manual work as well. If both refuges close, where do people go?

The signals are already here, and they are not encouraging.

The first concerns the wheel that has saved us so far: the birth of new tasks that only people can perform. Daron Acemoglu and Pascual Restrepo measured it in the United States. Until 1987, for every task machines took away from people, almost one new task appeared. After 1987, about half of one. The wheel still turns, but at half speed. And that was before artificial intelligence.

The second concerns those entering the world of work. In the United States, from late 2022 to the summer of 2026, the number of 22- to 25-year-olds employed in the occupations most exposed to AI fell by almost 13 percent, while their peers in less exposed jobs grew by 8 percent. In the last year the decline is 4.4 percent, more than twice that of their less exposed peers, and it is starting to reach those aged 26 to 30. Nobody fired them. They simply were not hired.

The third concerns who gets paid. For twenty years, all over the world, the share of the wealth produced that goes to those who work has been shrinking. This too began before AI. AI pushes in the same direction.

This book argues a thesis, and states it right away. Over the next twenty years artificial intelligence will make economies richer and, at the same time, reduce job opportunities across a large part of economic activity. New jobs will come, but not enough of them, not soon enough and not for the same people. The young will pay first, as the entrance door narrows without a sound. And unless the rules change, the value produced by machines will go mostly to those who own them.

To put this thesis to the test, a model of the world economy was built: three large regions, households divided by income, firms, a state that collects and spends. Four thousand possible futures, each calculated twice, with the new wave of AI and without it, starting from the same conditions. It is the most honest way there is to isolate the effect of the technology from everything else.

And here is what makes the result hard to ignore. The model was built on cautious assumptions, almost generous towards AI: it does not simulate machines able to do nearly everything, and it lets the exposure of activities grow by at most 60 percent. Even so, the picture is clear. The economy grows. But the poorer half of the

world consumes less than it would without AI, the share of wealth going to wages shrinks by almost six points, and in the advanced economies unemployment rises for a decade without returning to where it started.

This is what happens with cautious assumptions. The case in which AI and robots really do close both refuges is one the macroeconomic model excludes out of caution. Chapter 36 tackles it, with the book's forecast occupation by occupation, demand included: by 2046, for office clerks, about six jobs in ten fewer; for software developers, four in ten; for doctors and nurses, where demand for care grows, fewer than one in ten. Chapter 37 carries the same forecast from work to social tension, country by country.

The book is meant also for readers who have not studied economics. Each concept starts from the question it serves to answer and, when it helps, from an example with numbers. Tables, charts and formulas then lead into the technical analysis, and the more complex relationships are explained in the appendix, with the symbols and the steps.

Three rules apply to every page that follows. Every real figure has its source. The model's assumptions are listed once, at the end of chapter 6, and are not repeated on every page. The code is public: anyone can download it and redo the sums. The model has not been peer reviewed, and the open code exists precisely so that anyone can carry out that check.

One last thing, which should not be hidden. This book was also prepared with the help of artificial intelligence tools (Claude, ChatGPT and Grok) for parts of the text and the code. They are products of companies that sell AI and have every interest in telling a good story about it. That is one more reason not to trust the words and to check the numbers. The choices, the conclusions and the responsibility for what is written have not been delegated to any machine.

The title says that nobody can tell you for certain what AI will do to your job. This book tells you what is most likely, and why. Not knowing everything is no reason to do nothing. Those who wait for certainty arrive when the damage is already done.

Let us start from a question that seems trivial and is not: when we say an economy is growing, what exactly is growing? And for whom?

Chapter 1

Output, Employment and Income Distribution

In 2023, at an American customer support company, 5,179 agents get a new colleague. It has no desk, takes no holidays and never gets tired: it is an assistant based on artificial intelligence, which suggests answers while the agent is talking to the customer. The economists following the introduction of the tool measure everything. After a few months the agents solve on average 14 percent more problems per hour; the youngest and least experienced reach 34 percent more. The assistant has learned from the conversations of the best colleagues and passes that experience on to those who lack it.

For the company it is a gain in efficiency. For customers, shorter waits and better answers. For those who work there, the consequences depend on what happens next. Do orders increase? Do tasks change? Are hours cut? Are new people hired, or fewer? If a beginner with the assistant works almost like a veteran, and a veteran with the assistant does the work of one and a half, how many agents will be needed next year? The study does not say, because that was not its question. It is the question of this book: what happens to people when an activity produces value using less human work?

The technical improvement is only the starting point. A more efficient activity can coexist with households in difficulty, or bring more output and new job opportunities: productivity alone does not say which road prevails. Artificial intelligence and robotics change the relationship between output, work and the distribution of wealth in four ways: they increase what we are able to produce, they replace human activities, they make new services possible, they change who receives the income generated. This book measures how much each effect weighs, and at what pace.

To find your way you need three questions: how much is produced, how many people work, who gets the income. They are connected, but each looks at a different aspect. More output can

come with more employment, or with less. And even when total income grows, the benefits can be divided very differently among households. This chapter prepares the tools for reading these changes; the following ones present the data and the model, which covers the period from 2026 to 2046.

Two Economies with the Same GDP

Counting medical visits, completed files or delivered parcels helps to describe an activity. To compare a whole economy you need a measure that brings together very different goods and services. The most widely used is gross domestic product, GDP: the value added generated in a territory over a given period.

"Value added" is an expression to be taken literally. It serves to avoid counting the same thing twice along a supply chain: the miller sells flour to the baker, the baker sells bread, and GDP does not add up the price of the flour and the price of the bread, but what each one added. A firm that buys materials and sells a product contributes the value it adds, not the whole sale price. In the same way, a transport company does not produce the value of the goods it delivers, which may have been made elsewhere: its contribution is the transport service, net of the fuel and services used. Carrying an expensive product does not mean generating a large value added in transport.

Then the increase in output must be separated from the increase in prices. If the price of a service rises while the quantity stays the same, its value in euros grows, but no more services have been provided. When the book speaks of real growth it refers to measures that remove the effect of prices using price indices. The task becomes harder when quality changes or new services appear: what is a service worth, in real terms, that did not exist ten years ago?

GDP is governments' favorite thermometer, and it answers one question well: how much is produced. To understand how households live you need other information: incomes, distribution, job opportunities. And the difference is not academic. Imagine two small economies, identical at the start: both produce 100 and employ 100 people. Ten years later both produce 120, a fine 20 percent growth, the same headline. But in the first 110 people work, in the second 80.

Table 1 – Output and employment

Measure	At the start	Outcome A	Outcome B
Value added	100	120	120
Employed	100	110	80

Final output is identical, and the number of employed follows two opposite roads. In the first economy growth brought work. In the second, twenty people out of a hundred were left out, while the headlines spoke of recovery. Knowing that the economy has grown does not tell you which of the two columns you are in. Every time you hear that artificial intelligence will make GDP grow, this is exactly the question to ask.

Output alone does not say how many jobs there will be. And the employed do not tell everything either: eighty people can each work many more hours than a hundred, and those who work less may receive transfers or other income. To move from output to jobs you need to know productivity, demand for the products, hours, organization, skills.

Many apparent contradictions arise this way: by asking an indicator to answer a question it was not built for. A positive GDP figure can be true and leave the question of employment open. A negative figure on youth hiring can be true without the economy collapsing. The diagram that follows lines up the links the book will follow.

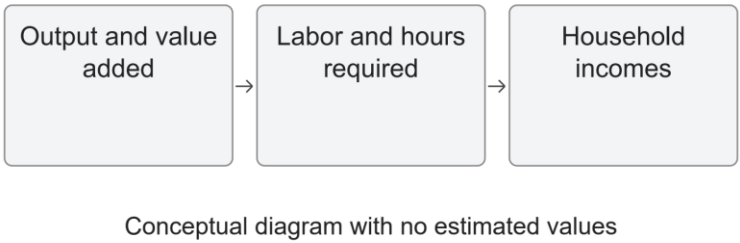


Figure 1 – Output, work and incomes are linked.

Who Gets the Value

For a household, growth becomes something concrete only through the money it receives and what it can buy. So the value produced must be followed to where it ends up.

Those who work receive a wage. Those who own capital receive profits, dividends or other income, depending on the form of the firm and its results. Part of all these incomes pays taxes and contributions, and the state returns it to households as services, pensions and benefits. These categories describe where income comes from, not separate groups of people: a worker can have invested savings, a pensioner can receive a pension and a rent. Speaking of labor and capital does not mean imagining two tribes.

The indicator that measures this division is the labor share: the part of total income that goes to those who work.

$$Labor\ share\ = \frac{Income\ attributed\ to\ labor}{Total\ income}$$

The income of a craftsman or a professional is partly labor and partly capital, and sources divide it in different ways: that is why, when data are compared, the definitions are checked.

An example shows why it matters. Total income goes from 100 to 120, and labor income from 60 to 66. At the start the labor share is 60 percent; at the end, 66 divided by 120, it is 55. Those who work receive more, six more, and yet their slice has shrunk, because the total grew faster. All the rest of the growth went elsewhere. Nobody stole anything: it is the way value is divided when machines do a larger part of the work.

Both pieces of information count, the one in euros and the one as a share. And for a household what also matters is its own income, what it has to pay and what it receives: a share calculated over the whole economy is not a pay slip.

A technology can make output grow and distribute the gain among wages, lower prices and capital income in very different proportions. The organization of firms, competition, bargaining and public policy change those proportions. And one basic fact already exists: over the last twenty years, even before generative AI, the world labor share has fallen. This book asks what happens if a new technology pushes in the same direction, and harder.

What Total Employment Hides

Every month the employment figure comes out. If it stays stable, it seems nothing is happening. But the number of employed is a snapshot: it does not say who came in and who went out to get there. To understand change you need the movements: who loses a job, who finds another, how long the passage lasts.

And you also have to look for those who could have come in and did not. A firm can keep all its experienced employees and simply reduce the hiring of beginners. The total does not move, nobody is fired, no headline. The change shows up first in the flows of entry rather than in the number of employed: the total looks stable while the prospects of those starting out worsen.

That is what is happening in the United States. From late 2022 to August 2026, among 22- to 25-year-olds, employment in the occupations least exposed to artificial intelligence grew by about 8 percent; in the most exposed it fell by almost 13. In the last year, from August 2025 to August 2026, the most exposed fell by 4.4 percent, the least exposed by 2. The difference does not come from mass layoffs, but from hiring that did not happen. Not all years weigh the same: 2022 and 2023 reflect post-pandemic over-hiring and higher interest rates, and the book counts them little. Chapter 3 explains how it is measured.

Those looking for a first job, moreover, face conditions different from those who have to find a new job after twenty years of work. Family responsibilities, education and the possibility of moving change the paths. And age alone does not say how exposed a person is to AI: what counts are the activities they perform, the occupation, the context.

That is why the book looks at the share of people who work, the share looking for work and the share who have stopped looking. A discouraged person disappears from the unemployed without having found a job. Chapter 3 explains the three measures and presents the first data by age.

The World That Does Not Exist

A firm adopts AI and, the following year, reduces staff. Is the technology to blame? Perhaps. Or perhaps it would have cut anyway, because customers were falling, or because it would have reorganized the work, or automated with other tools. Or perhaps, without AI, the cut would have been bigger, or it would have closed altogether. The sequence of events alone is not enough to give the technology the blame, or the credit.

To know what the technology did we would need to see a world that does not exist: the same, identical, without AI. Economists call it the counterfactual, that is, the comparison with what would have happened otherwise. In reality it cannot be observed. In a model it can.

That is what this book does. Every simulated future is calculated twice, in a pair of worlds with the same starting conditions: in one the new wave of artificial intelligence arrives, in the other it does not. The world without AI is not standing still: it keeps the same demography and the same ordinary technical progress. The only difference is the technology.

$$AI\ difference = Value\ with\ AI - Value\ without\ AI$$

When the book says that in 2036 the poorer half of the world consumes 1 percent less, it means: less than it would consume, in the same year, in the twin world without AI. It is the effect of AI calculated by the model. Chapter 18 explains how to check it against real data.

And the picture that emerges from four thousand comparisons is not what many would expect. There is no collapse. In 2036 world GDP is 1.3 percent higher, in 2046 3.2 percent higher: the economy grows more. But in 2046 the poorer half of the population consumes 3.3 percent less than it would without AI, the wage slice has shrunk by almost six points, and in the advanced economies unemployment rises for almost a decade and at the end of the horizon has still not returned to its starting point. A richer world, and not a fairer one.

Comparisons, moreover, must be on the same scale. A region next to the whole world, or one scenario next to an average of different scenarios, can look contradictory even when they are not. That is why, in comparisons about work, the book always states

territory, unit and sample. And the gap between labor demand and supply has a chart of its own, with a scale that does not squash it inside that of GDP.

Finally, in the book you will meet numbers of three kinds, and it pays to recognize them right away.

Table 2 – Three kinds of information

Information	What it represents
Observed data	A measure with a source, territory, period and definition
Teaching example	Numbers chosen to explain a mechanism
Simulation	A result conditional on the equations and assumptions

Observed data are measurements of the real world, with their source: the 34 percent of call center beginners, the fall of young people in exposed occupations. Examples, like the two economies in table 1, serve to show a mechanism. Simulation results are the consequences of the model's equations and assumptions, listed once at the end of chapter 6. The assumptions are written down and the code is open: anyone who disagrees can change them and redo the sums.

The book also uses a second tool: a module that follows people by age, that is, those entering for the first time, those losing their job, those becoming discouraged. It receives labor demand from the main model and simulates entries, exits and discouragement. Chapter 26 closes the circle: it feeds the module's unemployment back into the wages and benefits of the main model, and shows how much the results change.

To follow the book you do not need to know all the tools in advance. In the text only the simple relationships that help to understand an example appear; the more complicated formulas are collected in the appendix, with questions, an explanation of the symbols and examples, and each cross-reference points to the

right section. In front of a chart it pays to look for three things: the question, the axes, the term of comparison. And chapter 12 explains how to redo the calculations, so that the discussion can go all the way to the numbers.

The opening question can now be asked more precisely. To understand what happens when less human work is needed you have to measure output, follow employment and reconstruct the distribution of incomes. The next chapter starts from what comes before all these measures: work itself. A job is not a single block, but a set of activities, some of which a machine can already perform. The useful question is not whether AI will replace you, but which of your activities it will replace, and what changes when it is really used.

Chapter 2

Tasks, Exposure and AI Adoption

Try to describe any job, yours or that of someone you know, with a single verb. It cannot be done. People in accounting collect documents, verify information, apply rules, spot anomalies, talk to clients. People who write software interpret requests, write code, check results, decide how to fix a problem. Even work that looks repetitive contains situations that call for attention and judgment.

To understand what artificial intelligence will do to jobs, this is where you have to start. Automating one phase can reduce the hours needed, change responsibilities, or make it worthwhile to expand the business. Before talking about jobs you have to reconstruct what work is made of, and distinguish what a technology can do from what an organization decides to entrust to it.

A Job Made of Tasks

In research, the activities that make up an occupation are called tasks. The distinction between task and occupation is the starting point of this book's model, and of almost all economic research on automation over the past twenty years.

A technology can perform some tasks without being able to do the whole job. It can find information in a second, prepare a first draft, run routine checks, and leave a person the review and the final decision. In other cases it replaces an activity completely enough to reduce the working time needed. And even when all the visible activities seem automatable, questions of organization remain. Who answers for a mistake? Who changes the procedure when conditions change? What new activities arise with the new system? A machine does not always need a person beside it; but a technical capability does not immediately become operational autonomy.

Take work in a port, or in a large warehouse. Very different tasks coexist within the same job, and each poses a different question to anyone thinking of automating it.

Table 3 – Activities within an occupation

Task	Possible use of the technology	Question to check
Documents	Information extraction	Reliability of the results
Routes	Optimization	Handling the unexpected
Handling of goods	Autonomous vehicles and cranes	Operating conditions
Anomalies	Analysis and reporting	Decision and responsibility
Emergencies	Information support	Coordination and limits

The table lists the checks to make before turning a technical capability into an estimate of jobs.

That a technology can do something, moreover, does not mean it is used. This is the step from exposure to adoption, and confusing the two is the most common mistake in newspaper headlines. Exposure says that an occupation's activities can be touched by the technology. Adoption says that the technology is actually used. They are two different quantities, each with its own denominator.

An occupation can be highly exposed and work in firms that adopt slowly. A system can work very well and be too expensive for a small business. Usable data may be missing, or integration with the rest of the work, or training, or authorizations. Or adoption suddenly accelerates when cheap tools that are easy to fit into everyday work arrive. What counts is the result in a real organization, not the laboratory test: quality, reliability, total cost. And checking costs something too: a draft written in a second saves little if it then takes two hours to verify.

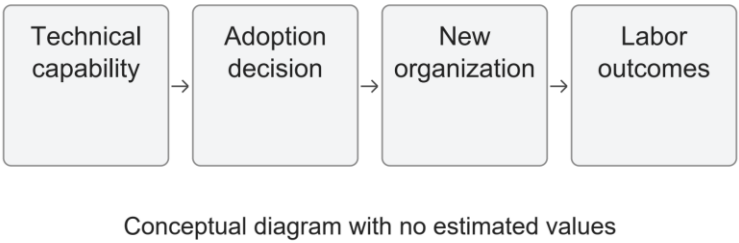


Figure 2 – From technical capability to change in work. The steps may slow down or stop; the diagram does not indicate a probability and does not imply that the final outcome is a loss of employment.

In the model, the new wave of adoption begins in 2026, on top of the automation firms already use. And technology does not enter all firms at once. Like almost every innovation of the past, it starts slowly, accelerates and then slows near its limit: an S-curve. Exposure, by contrast, says how much of the activities can be touched, and it grows even before firms adopt the tools; in the model it does not exceed 95 percent. The shape of the curve is the assumption that weighs most on all the results, as chapter 16 shows. Section A1 of the appendix explains the calculation.

Forty Percent of What?

You may have read that artificial intelligence concerns 40 percent of the world's jobs. That is the International Monetary Fund's estimate, in its 2024 report Gen-AI: Artificial Intelligence and the Future of Work. The International Labour Organization, with the Polish institute NASK, in its 2025 report Generative AI and Jobs: A Refined Global Index of Occupational Exposure, looks instead at the highest band of exposure to generative AI and arrives at 3.3 percent of world employment.

The two figures do not contradict each other, because they answer different questions. The first measures a wide perimeter: all the jobs in which some of the tasks can be touched by the technology. The second measures a core: the most exposed jobs, in a classification built by another study, with another definition of technology and another threshold. And neither of them says how many jobs will disappear.

Table 4 – Measures of AI exposure

Measure	Meaning
Exposure	Technical possibility according to the source's method
Adoption	Actual use of the technology
Substitution	Transfer of tasks from labor to technology
Employment	People employed within the defined scope

That is why the book accompanies every percentage with the question asked by the source. Those who want to frighten you will use the 40; those who want to reassure you, the 3.3. Neither is telling you what will happen to your job.

And age does not replace a measure of activities. To say that the young are more exposed you need to know which occupations they hold, which tasks, with which classification.

The Simplest Formula and Its Traps

The simplest measure of labor productivity relates how much is produced to the hours used to produce it.

$$p = \frac{Q}{H}$$

Q is the quantity of a well-defined product or service, H the hours of work, p how much is produced in an hour. You can count completed applications or pieces made, knowing that different

units require different amounts of time. To compare whole economies, you need measures of value added and prices instead.

Suppose an office completes a hundred applications in a hundred hours: productivity is one application per hour. A technology arrives, and the same hundred applications take eighty hours. Productivity rises to 1.25. The increase is 25 percent, while the hours needed fall by 20. The two percentages differ because they start from different bases. So far, good news. But how many people are needed?

Table 5 – Productivity and labor required

Situation	Quantity Q	Hours H	Q / H
Before	100	100	1.00
After, same Q	100	80	1.25
After, Q +25%	125	100	1.25
After, Q +50%	150	120	1.25

It depends on how many applications come in. Productivity stays at 1.25 in every case after adoption, but the hours needed change with the quantity required. If applications stay at a hundred, eighty hours are enough: a fifth of the work is no longer needed. If they become 125, a hundred hours are needed again. If they become 150, a hundred and twenty are needed, more than before. The relationship can be rewritten like this:

$$H = \frac{Q}{p}$$

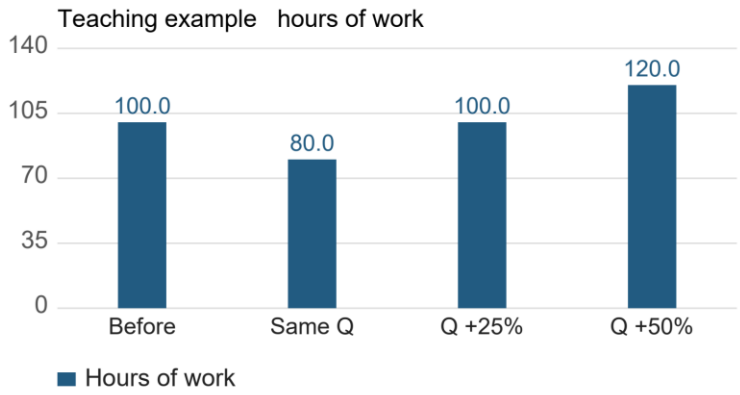


Figure 3 – Hours required in an example with higher productivity. The columns come from the teaching table and do not represent observed data or a forecast about AI.

The formula contains no wages, skills or investment. But it contains the heart of the debate on AI. Optimists bet on the last column: prices will fall, demand will grow, more work will be needed. Pessimists bet on the first: demand will not grow enough. The answer depends on how much more paperwork, advice, translation and software the world will want when they cost less. A rise in productivity, on its own, does not say what happens to jobs.

In the model, productive capacity grows through an underlying trend and three contributions from AI: savings on automated tasks, help to people in the tasks that remain, and the creation of new tasks. The weights of the three contributions are among the assumptions listed at the end of chapter 6. And the capacity to produce more must then find someone to buy it. Section A2 of the appendix shows the full calculation.

Hours, moreover, are not jobs. Two firms can use the same hours with a different number of people: one all full-time, the other with more people on shorter hours. A fall in the hours needed can become less overtime, contracts not renewed, fewer hires, shorter hours, or layoffs. Even when the job remains, the work

can change: tasks, autonomy, pace, pay, responsibility. A job kept is not necessarily the same job as before.

And to talk about employment you have to count people, not contracts. A person can have several contracts in the same year; a firm can hire someone who already works elsewhere. That is why administrative hiring records and surveys of people answer partly different questions. Chapter 3 starts from exactly this difference.

A Race Between Two Forces

Think of a transport service that becomes cheaper thanks to automated driving. The lower cost increases demand and recovers part of the labor saved. How much is recovered depends on volumes and on the activities required: not all drivers will become robot supervisors.

The model distinguishes two forces. On one side there is displacement: some activities pass from people to capital, which includes technologies. On the other there is the creation of new tasks, or the recovery of old ones, in which human labor once again contributes to production. In the theoretical framework the book uses, that of Acemoglu and Restrepo, this second force is called reinstatement. It is the force that, throughout history, has brought work back after every wave of machines.

New tasks do not always correspond to new, recognizable occupations: they often enter jobs that already exist. And their appearance does not say how many jobs will be created: that requires the demand for the activity and the time needed to perform it. How skills are distributed matters too: a job can stay vacant while a person looks for work on the same street, because the two do not meet.

And timing matters. The replaced activity disappears quickly. The new one appears years later. Meanwhile a family has to pay

the rent, the groceries, perhaps a training course. A positive balance after twenty years does not tell you what happens in the twenty years in between.

In the model, the human labor content falls when tasks are replaced and rises when new ones appear, and the two movements fall in different years. The world without AI keeps the underlying changes. Bargaining can return part of the lost income to labor, but without creating jobs. Section A3 of the appendix explains the relationships between tasks and the labor share.

Retraining, finally, is not instantaneous. A person may need a course, practical experience, someone to give them a first chance. They may face economic, geographical or family obstacles. A system that creates new tasks does not guarantee that job seekers have the skills required or can reach the jobs.

That is why the book adds to the main model a module that follows people by age: first entry, re-employment, discouragement. The advanced economies start from the situation of the European Union. Among its case studies there is one with slow retraining and one with persistent displacement: they show what happens when new tasks do not offset displaced work quickly enough.

To check what is really happening, you have to connect activities and people: you need data on adoption, on the composition of occupations, on entries and exits. And to say that AI is the cause, you have to separate it from the business cycle, from demand in a sector, from other organizational changes. Comparing more and less exposed activities helps, because the two groups live through the same cycle; looking at the previous period and at the composition of the sample helps rule out alternative explanations. That is why every chart in the book says who is counted, in which period, with which unit. An index curve is not an unemployment rate, and a figure from a sample is not a country's total.

There is also a consequence few people see. The tasks easiest to automate are often those given to beginners: the first files, simple searches, drafts to be corrected. They were dull, but they were also the way one became an expert: making mistakes on small things, watching those who knew how, earning a little responsibility at a time. If that path disappears, we need to understand where the experience required for later tasks will be built.

The next chapter follows exactly this door. It distinguishes the people who already work from the movements that bring others into work, presents the first observed data and asks a question that aggregate numbers tend to hide: how many opportunities to begin remain open?

Chapter 3

Participation, Unemployment and Discouragement

A firm can keep all its experienced staff and simply stop hiring beginners. For a few years the number of employees barely changes. Nobody is fired, no union protests, no newspaper reports it. But the chances of starting a career shrink. This is one of the reasons why, to understand the future of work, it is not enough to look at how many people work: you have to look at who manages to get in. This chapter defines these movements and the indicators that make it possible to follow them, and presents the first observed data.

Getting In and Learning

Many trades are learned step by step: first relatively simple activities, then watching colleagues, having mistakes corrected, taking on growing responsibility. Part is learned at school or university; another part is learned only by taking part in real situations. For centuries this was the model of apprenticeship. In England, the Statute of Artificers of 1563 set the length of apprenticeship for many trades at seven years: seven years beside a master before you could work on your own. The forms have changed, but the idea has remained: experience is built by doing, starting with small things.

When a technology performs the initial activities, the organization gains efficiency. But it has to find another way to prepare those who will one day handle the complex ones. AI can also help people learn: giving immediate feedback, making accessible activities that used to require years of experience.

The two things can coexist. A young person who has already been hired may receive tools that make them more capable, while other young people find fewer positions to start from. The productivity of those who get in and the probability of getting in are different questions, and they require different measures. A

learning path interrupted today weighs on tomorrow's skills; measuring it requires data on access, initial tasks, training and careers.

To see it in the numbers, you have to distinguish two things. The number of people employed is a snapshot at a moment: economists call it a stock. The movements that change the snapshot are called flows. If you confuse the two levels, you can miss a transformation that is under way.

$$E_{t+1} = E_t + I_t - X_t$$

E is employment, I the entries into work during the period, X the exits. The relationship holds if the perimeter is consistent and includes all movements: if the territory or the observed population changes, that must be taken into account. X is used so that exits are not confused with U, the unemployed.

Table 6 – Entries into and exits from employment

Movement	First year	Second year
Initial employment	1,000	1,000
Entries	100	60
Exits	100	60
Final employment	1,000	1,000

In the example, employment stays at a thousand in both years. But entries fall from a hundred to sixty. Exits fall too, so the total does not move. The entrance door has narrowed by 40 percent without the employment snapshot showing anything. That is why flows must be measured. The opposite can also happen: an activity loses employment and still hires many people, because exits are even larger. The balance and the movements tell complementary stories.

Flows can also count repeated episodes: a person can enter, leave and come back in the same year. Depending on the source, people or events are counted. An administrative series of contracts cannot be compared directly with a model that allows only one hire per person in a period.

And not all hires are the same. A firm can hire someone in their first job, someone returning to work after a spell of unemployment, or someone moving directly from another job. For the firm they are all new hires. For the labor market they are different movements.

Table 7 – Hiring and first entry

Movement	Previous situation
First entry	No work experience within the defined scope
Re-employment	Previous job followed by a period without work
Direct job change	Move from another job
Return from inactivity	The person was not working and not actively looking

First entry requires knowing a person's history. Sometimes we know they have never worked; sometimes only that they do not appear as employed in the available records. Re-employment requires knowing a previous job and the period without work that followed. The quality of the data determines how well the paths can be distinguished. Age helps identify groups but does not replace history: a twenty-four-year-old may have been working for years, an older person may be looking for their first paid job. Charts by age and charts by type of entry answer different questions.

The module added to the model distinguishes those looking for their first job from those looking after previous experience, and

gives them different entry probabilities, set among the model's assumptions.

Who Counts as Unemployed

Labor statistics divide the population into employed, unemployed and inactive, according to precise criteria. The labor force is the sum of the employed and the unemployed. And not having a job is not enough to be unemployed: you also have to be looking for one and be available, according to rules set by the source.

The inactive are a galaxy: people who study, people who look after their family, people who are retired, people in other situations. They are not all discouraged. Discouragement is a specific phenomenon: stopping the search because opportunities seem too few. Recognizing it in the data requires additional information.

Three indicators are used more than any others. P is the population of the age group considered, E the employed, U the unemployed; the labor force is E plus U.

$$Employment\ rate\ = \frac{E}{P} \times 100$$

$$Unemployment\ rate\ = \frac{U}{E\ +\ U} \times 100$$

$$Participation\ rate\ = \frac{E\ +\ U}{P} \times 100$$

The employment rate says how many people work out of the whole population of the group. The unemployment rate says what part of the labor force is looking without finding. The participation rate says how much of the population is in the labor force. The denominators are different. That is why employment and unemployment are not added up as if they had to make a hundred, and 20 percent in one indicator does not mean the same thing as 20 percent in another.

There is a way for unemployment to fall without a single job being created. Take a hundred people: sixty work, ten are unemployed, thirty inactive. The labor force is seventy. The employment rate is 60 percent, the unemployment rate about 14.3. Now five of the ten unemployed people stop looking and become inactive. Employment stays at sixty. The labor force falls to sixty-five. The unemployment rate drops to about 7.7 percent. The employment rate stays at 60.

Table 8 – Discouragement and denominators

Measure	Before	After
Population	100	100
Employed	60	60
Unemployed	10	5
Inactive	30	35
Employment rate	60.0%	60.0%
Unemployment rate	14.3%	7.7%
Participation	70.0%	65.0%

The headline would say: unemployment halved. But not a single job has been created. Five people have simply stopped hoping. That is why employment, participation and the reasons for inactivity have to be looked at together.

To avoid losing these people, the model's extension keeps the discouraged in a measure of potential supply: everyone who would work, including those who have stopped looking. It is a measure of the model, different from official indicators of labor underutilization and from young people not in education or employment.

The First Observed Snapshot

Let us move from examples to data. The European Labour Force Survey, published by Eurostat, makes it possible to compare Italy with the twenty-seven countries of the European Union. For this book, annual data from 2019 to 2025 were archived and four age groups built: 15–24, 25–34, 35–49 and 50–64. To combine several age groups, employment and population are added up and then the ratio is calculated, instead of averaging the rates: larger groups must weigh more. The source's flags are kept, and the possible methodological change of 2021 is taken into account.

Table 9 – Observed data for 2025, rounded. The employment rate has the population of the age group as its denominator; the unemployment rate, the labor force of the age group. Source: Eurostat, series lfsa_pganws, aggregation of the five-year age groups archived in the project.

Age	IT empl. %	IT unempl. %	EU empl. %	EU unempl. %
15–24	17.9	20.6	34.5	15.2
25–34	68.5	9.0	79.6	7.2
35–49	77.5	5.5	84.4	4.8
50–64	66.5	3.8	71.9	4.4

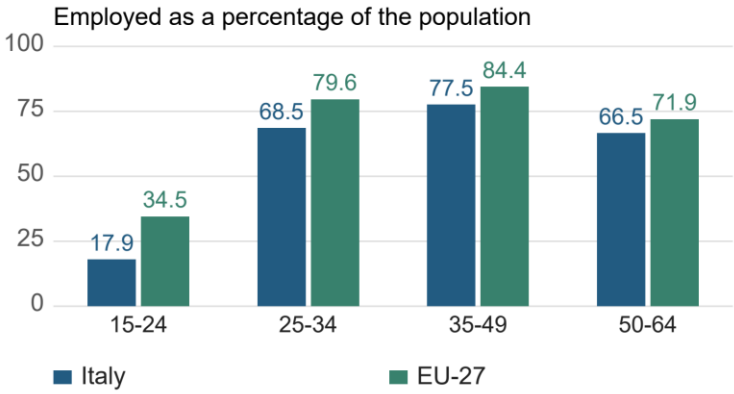


Figure 4 – Share of people employed in 2025 by age group. All occupations, Italy and EU-27. Source: Eurostat.

In Italy, in 2025, fewer than one in five young people aged 15 to 24 work: 17.9 percent, against 34.5 in the European average. Among those looking, one in five does not find a job: youth unemployment is at 20.6 percent, against 15.2 in Europe. In the older age groups too, Italy remains below Europe in employment: 68.5 against 79.6 between 25 and 34, 77.5 against 84.4 between 35 and 49, 66.5 against 71.9 between 50 and 64.

The figure for the very young is not pure exclusion: many are studying, and differences between countries also reflect different study paths, work during studies and the structure of the economy. These differences have roots much older than AI. But the snapshot does say one thing: in Italy the entrance door was already narrow before artificial intelligence arrived.

A Signal Among Young Americans

To understand the effect of a technology you need a comparison between similar activities, taking into account customers, products and changes already under way. For AI, the most closely watched figure today comes from the United States. The Canaries dash-

board, built by Stanford researchers with payroll data from ADP, one of the largest American payroll processing companies, follows employment in groups of occupations more or less exposed to AI. The name comes from the canaries miners took underground: when the canary stopped singing, there was gas in the air. The sample is that of ADP's client firms: it is American, and it does not cover the whole market.

The source's age groups are 22–25, 26–30, 31–34, 35–40, 41–49 and 50 or older. They differ from the Eurostat groups, and the book leaves them as they are. To read the chart, all you need is the index with base one hundred.

$$Index_{\mathrm{t}} = \frac{Value_{\mathrm{t}}}{Initial\ value} \times 100$$

If a group goes from two hundred to one hundred and eighty employed, its index falls from a hundred to ninety. If another goes from two thousand to two thousand two hundred, its index rises to a hundred and ten. Indices compare relative changes, not numbers of people, and the distance between two indices is not an unemployment rate.

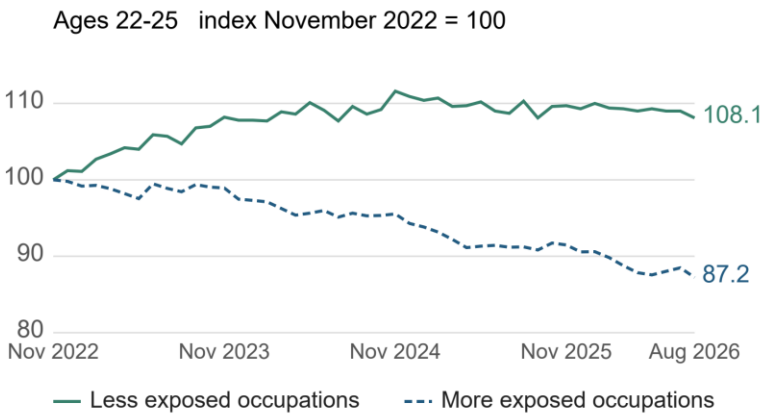


Figure 5 – Employment of workers aged 22 to 25 in the Stanford/ADP sample, less and more exposed occupations. Base November 2022 = 100. Data version of September 17, 2026, observations through August 2026.

In August 2026, with November 2022 set at one hundred, the index for 22- to 25-year-olds is about 108.1 in the least exposed occupations and 87.2 in the most exposed. The former have grown by about 8.1 percent, the latter have fallen by about 12.8. The gap is about 20.9 index points: not points of unemployment, nor a percentage of jobs lost across the economy. In less than four years, among peers of the same age in the same market, the two paths have split by twenty points.

Not all these years say the same thing about AI. 2022 and 2023 reflect post-pandemic over-hiring and higher interest rates; 2025 and 2026 are the years in which generative AI really enters firms. That is why the book gives each year a weight: 100 percent to 2026 and 2025, 50 to 2024, 5 to 2023, 1 to 2022. Every monthly change counts with the weight of its year. With these weights, 22- to 25-year-olds in the most exposed occupations lose 4.2 percent a year, those in the least exposed 0.3. Between 26 and 30, minus 2.2 against plus 0.3. In the last year, from August 2025 to August

2026, the decline of the most exposed is 4.4 percent against 2 for the least exposed, and between 26 and 30 it is 3.1 against plus 0.1.

The authors of the research, analyzing the flows, show that the divergence arises mainly from fewer hires, not from more exits: chapter 36 returns to this. The series kept in the project starts in September 2021, before the base of the chart.

Learning on the job connects today's work with tomorrow's skills. Shadowing an experienced person costs something right away, but builds capabilities that already automated activities no longer teach. If entry-level positions shrink, the training of future experts can change. One firm can respond with new paths, simulations, mentoring and AI tools; another can postpone the problem. The outcome cannot be read from the number of hires in a single year: it has to be observed over time.

In what follows, the book measures first entry, job loss and reemployment separately in the simulations. The Eurostat transition data have no ages, and for Italy youth flows require administrative sources and microdata, which the book has not yet used.

In this way the question becomes testable: how many people can start, which skills do they learn, where do they find their next job? Employment figures give part of the answer; paths give another. And the difficulties can be two, even at once: people who need work but lack the skills or the means to reach the jobs, or too few jobs for everyone who is looking. The next chapter introduces labor demand and supply and the difficulty of matching them, with charts that keep the world and the advanced economies apart, as well as those actively looking and those who would work if they could.

Chapter 4

Labor Demand and Supply

A hundred people looking for work and a hundred vacancies do not necessarily make a hundred hires. The right skills may be missing, or the transport links, or hours compatible with the lives of those looking. If instead there are only eighty jobs, matching people and jobs better helps, but it does not change the fact that there are not enough jobs. They are two different problems, and to separate them you first need clear definitions of demand, supply and units of measurement.

Who Supplies and Who Demands

In economics, the words demand and supply seem upside down compared with everyday experience. A person "asks" for a job, but in fact supplies their labor. A firm "offers" a job, but in fact demands labor to produce goods and services. In this book, labor demand is the need of firms and other organizations; labor supply is the people, hours and skills available.

This is not pedantry: it prevents misreading the charts. If labor supply grows, it does not mean firms are posting more vacancies. If demand falls, it does not mean people want to work less. Before interpreting a curve, you need to know whom it represents.

And you need to know what it counts. An organization may need a certain number of hours, a certain number of people, or very specific skills. In the model's equations, demand is measured in normalized labor units; in the charts it becomes an index.

Table 10 – Labor demand and supply

Actor	Action	Measure
People	Supply labor	Hours, people, skills
Organizations	Demand labor	Needs for production
Model	Compares the quantities	Normalized units

The previous chapter introduced the active labor force. The age extension also follows those who have stopped looking, and keeps them in potential supply. That way a difficulty shows up that would become invisible in unemployment alone. Chapters 7–9 describe who belongs to it.

In its simplest form, the gap between supply L and demand D is this:

$$F = L - D$$

A positive value means there are more people available than work required; a negative value, the opposite. The formula measures quantity; the difficulty of matching skills and jobs comes afterwards.

In the original macroeconomic model, employment depends on demand, with a ceiling linked to supply, and there is a structural unemployment rate set in the calibration. So even the world without the new wave of AI has some gap between demand and supply, which should not be attributed to artificial intelligence.

To separate the two, each simulation compares demand with AI and without it. If supply is the same in both worlds, the additional gap is simply the difference between the two demands:

$$\Delta F = D_0 - DAI$$

ΔF is the additional gap, D_0 demand without AI, DAI demand with AI. In general, a wider gap can come from lower demand, higher supply, or both. And in the age extension active supply also changes, through discouragement: there the book keeps the movements of demand and supply separate.

The Chart That Hid the Gap

The first versions of the charts in this project had a flaw, and it is worth telling because it is a very common mistake. They put the advanced economies next to the whole world, and compared one scenario with the median of all simulations. They were not views of the same comparison. The revision separated territory, sample and scale, so that differences in composition would not look like economic contradictions.

Part of the difference came from geography: the labor force could shrink in the advanced economies and grow in the world as a whole. Another part came from the selection of simulations, which made the comparison between the two world pages unclear. A comparison attributed to AI must use the same sample.

And there was the scale. When GDP and other growing indicators take up almost the whole height of the chart, the distance between the labor curves becomes visually tiny. That does not mean it is zero. In the final indices, for example, world supply reached about 115 and demand with AI about 106. The difference was there, but next to a GDP close to 178 the eye could not see it. A chart can tell the truth and hide it at the same time.

For the book, the charts were rebuilt directly from the code. Both versions, with and without AI, use all the same four thousand draws, without selecting a scenario in a single curve. Demand and supply have a space of their own. The scale is consistent within each comparison, and the base of the indices is stated.

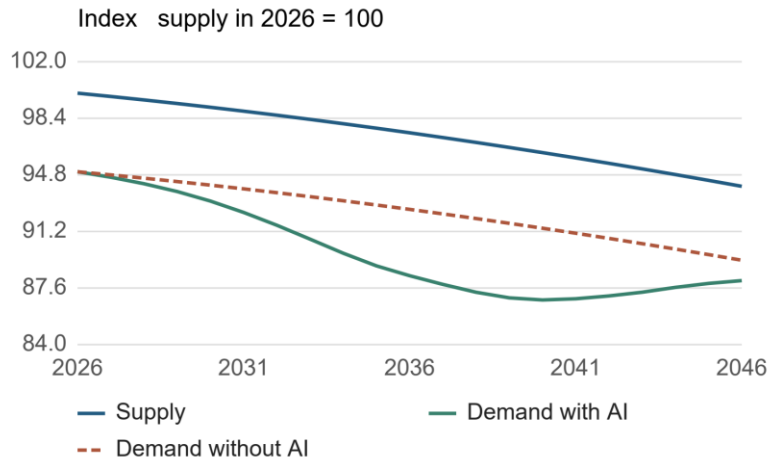


Figure 6 – Advanced economies, original macroeconomic model. Medians of the same 4,000 paired simulations, 2026–2046. Base 100 = labor supply in 2026. Initial demand is below one hundred because the model includes structural unemployment. Simulated results, not observed data. Model simulations.

In the advanced economies supply falls, following the demographic assumptions: fewer young people, more retirees. Demand without AI keeps the gap associated with structural unemployment. Demand with AI moves away during the spread of automation, and in the second half of the period it recovers only part of the distance.

Two different gaps need to be read here. The first is between supply and demand with AI: it is the overall gap in the model. The second is between demand without AI and demand with AI: it is the change due to the new wave of technology in the simulated pair. They are not the same measure. The differences are expressed in percentage points relative to the initial structural unemployment rate of the advanced economies, which the model sets at 5 percent. Chapter 14 presents the full picture.

So the fall in supply does not solve the problem on its own. "There will be a shortage of workers, so AI is arriving at just the

right time" is an argument often heard. It works only if demand falls more slowly than supply and if new tasks, output and wages recover enough. In the model this does not happen.

The world, moreover, does not have a single trajectory. The model divides it into three large regions: advanced economies, emerging and middle-income economies, low-income countries. Within each region there are huge differences: the regions represent common dynamics, not individual countries.

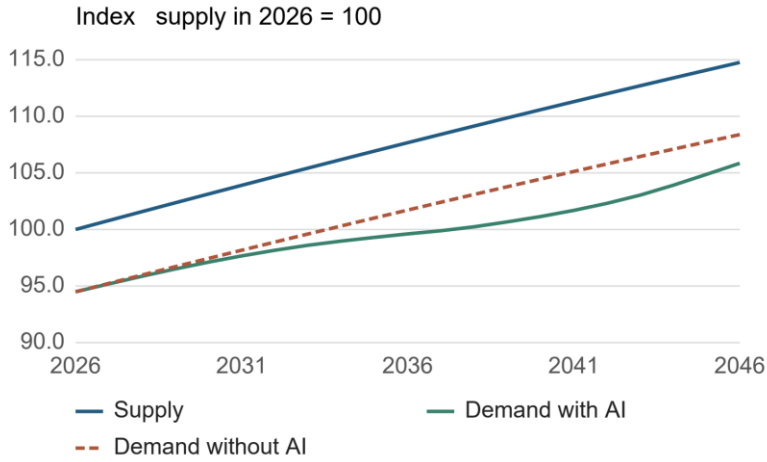


Figure 7 – World aggregate rebuilt with the regional weights of the original model. Same 4,000 draws for the curves with and without AI. Base 100 = world supply in 2026. The weights are regional approximations from the calibration, not an official series of the world labor force. Model simulations.

In the world as a whole supply grows, because the regions with rising populations offset the decline in the advanced economies. Demand can grow too and still stay below supply. And it can grow with AI and still stay below the demand there would have been without AI.

This distinction recurs throughout the book: improving compared with the start does not mean improving compared with what would have happened otherwise. Growing labor demand can

coexist with greater difficulty than in the world without AI. To measure the effect of AI, the book first calculates the differences within each pair of worlds, then their medians; chapter 6 explains why. The figures in this chapter show levels to make reading easier.

When Jobs Exist but Stay Empty

An overall gap does not tell the whole story. In 2010 the Nobel Prize in economics went to Peter Diamond, Dale Mortensen and Christopher Pissarides precisely for explaining why, in a labor market, vacancies and unemployed people can coexist for a long time: searching for each other and finding each other takes time and effort. The United States saw this spectacularly in 2022, when there were about two vacancies for every unemployed person while millions of people were looking for work.

Imagine an organization looking for ten technicians to maintain robots, and ten people with experience in manual handling of goods looking for work. Ten jobs and ten candidates, but the match is not guaranteed. The skills may be missing, or the willingness to move, or acceptable working conditions. Economic research calls the meeting of demand and supply matching, and the misalignment between the characteristics of jobs and those of people mismatch. A person can look for work in the same city where a firm is looking for different skills, and they may never meet.

The original model represented adjustment in aggregate form. The age extension adds probabilities of entry and re-employment, and allows some jobs to remain vacant.

People looking for their first job and people who have to change trade face different obstacles, and the age module makes it possible to follow them separately. There are four age groups: 15–

24, 25–34, 35–49 and 50–64, and people age over the course of the simulation, moving from one group to the next. To distribute demand across groups, the model uses accessibility weights, listed among the assumptions in chapter 6.

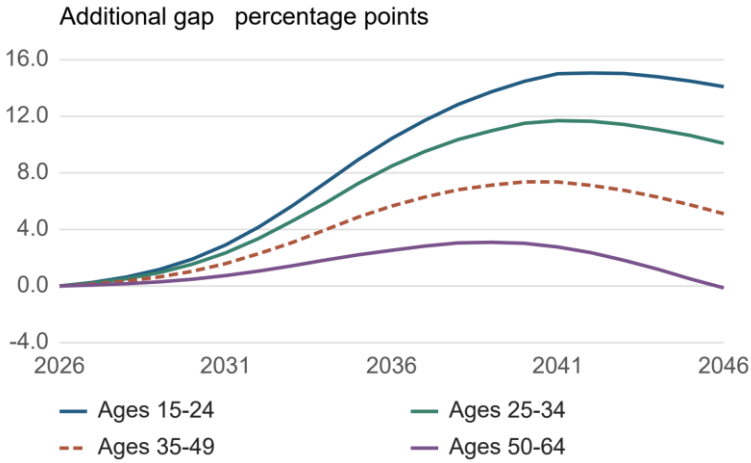


Figure 8 – Advanced economies, age module, slow-retraining scenario. Difference between the potential gap with AI and without AI, calculated within each pair and then summarized by the median. Percentage points of contemporaneous potential supply. The four age groups keep the model's classification. Retraining parameters assumed, not estimated.

This measure keeps the discouraged within potential supply: if it rises, the gap worsens compared with the pair without AI. The world without AI in the pair has the same slow market, so AI is not blamed for the difficulty that would exist anyway. The curve shows why the speed of retraining is one of the most important things to know.

The Price of Labor Does Not Solve Everything

That leaves wages. Cheaper labor can make it worthwhile to use more people in some activities. This is the classic answer: if there is unemployment, wages fall and firms hire again. But how well it

works depends on the demand for the service, the technology available and skills.

The model was revised precisely on this point. The first version let labor demand close the gap with supply too easily. The one used in the book derives demand from output and tasks, and gives it limited sensitivity to wages. To see what this means, take an elasticity of 0.3: if the wage falls by 1 percent, holding output and everything else fixed, the quantity of labor demanded rises by about 0.3 percent. Section A4 of the appendix explains the relationship.

And on its own, this relationship does not see the consequences of wages for consumption. If a household earns less, it spends less. What lowers costs for one firm can lower revenues for another. In the model, a lower wage can make hiring worthwhile, but the recovery is limited. Firms demand labor according to output, human tasks and the relative cost of labor. Employment cannot exceed available supply, net of structural unemployment. Wages respond to productivity and to the state of the labor market. And bargaining changes the division of income without magically creating jobs. Section A4 keeps these steps distinct.

The gap between demand and supply describes a problem of quantity. The accessibility of jobs adds skills, places and timing. To understand the consequences for people, you also have to follow the money: what is available while searching, and what arrives after being hired. The next chapter follows wages, profits, transfers and consumption, and shows how the same increase in productivity can become very different experiences for different households.

Chapter 5

Labor Income, Capital Income and Consumption

In January 1914 Henry Ford announced that he would pay his workers five dollars a day, about twice the going wage. Ever since, the story has been told like this: Ford wanted his workers to be able to buy the cars they built. Historians are more cautious. The main reason was something else: on the assembly line the work was so hard and repetitive that workers left in droves, and in 1913 Ford had to hire, in a single year, several times its entire workforce just to replace those who quit. The high wage was meant to keep them.

But the legendary version has survived for a reason: it contains a true idea. A wage is not only a cost for whoever pays it. It is also the income with which someone will buy something. This chapter follows the value produced by a more efficient activity to where it ends up: wages, property income, lower prices, investment. The path depends on who receives the resources and how they use them, and to understand it you have to move from output to household income.

A Household and Its Income

A household can receive wages, self-employment income, pensions, benefits, income from assets. And it pays taxes and contributions. What remains to spend and save is called disposable income. To reason about it, a simple identity is enough:

$$Yd = Income - T$$

T is the taxes and contributions. Income is the items listed, counted once: a pension or a benefit is not added twice.

Disposable income in euros, however, is not enough. You need to know what it buys. The same sum allows different choices if the cost of housing, food, transport or other services changes. Talking about real income means keeping money and prices together. And

an average price does not describe every experience: households that spend differently are affected differently. A service made cheap by AI can be vital for one person and irrelevant for another.

On the other side there is the firm, which does not automatically turn every hour saved into profit. It can invest, pay for technology services, lower prices. It can sell more, or lose customers. To understand what happens, you need to know all the costs and all the revenues.

The model uses an aggregate representation that distinguishes labor income from capital income. It serves to study how distribution changes when human tasks decline. But the labor share does not on its own determine each person's wage: how much is produced and how many people work also matter. In simple form, the wage bill is written like this:

$$W = w \times E$$

W is the wage bill, w the average wage per unit of labor, E employment. If both w and E change, the wage bill can fall in two ways. And there is a trap in averages: the average wage can rise when the lowest-paid workers leave, without those who remain getting a single euro more. Composition is part of the reading.

For capital, you have to distinguish wealth from income. Wealth is what is owned at a moment: houses, plants, shares. Income is what that wealth yields each year. The value of a plant is not added to annual income as if they were the same thing. And a share of wealth is not a share of income, even though the two are linked.

The Pie Grows, the Slices Change

GDP can grow while the income of some households falls compared with the world without AI. This is one of the central mech-

anisms of this book, and it becomes clear when you separate the growth of the pie from how it is cut.

Take an economy that produces value added of a hundred: sixty go to labor, forty to capital. After a technological change it produces a hundred and twenty. In a first outcome labor receives more in euros but a smaller share. In a second outcome labor gets fifty-five and capital sixty-five: output has grown by 20 percent, and total labor income has fallen.

Table 11 – Income distribution and consumption

Measure	Before	Outcome A	Outcome B
Value added	100	120	120
Labor income	60	66	55
Capital income	40	54	65
Labor share	60%	55%	45.8%

Final output is identical in the two outcomes. In the first, labor gains in absolute value and loses share; in the second it loses in absolute value too. To arrive at household incomes you then have to add taxes, transfers, property and the distribution of wages.

And even widespread growth does not reach every person. A model that groups millions of people loses the differences within groups. The median of a group does not describe a dismissal, an indebted family, a person who finds a better job.

In the model, household income comes from labor, capital and transfers, net of taxes. Distribution depends both on wages and on ownership. Each group consumes part of its income; what it saves, together with the profits firms retain on behalf of their owners, adds to wealth, and so changes the distribution in later years. It is a self-reinforcing mechanism: those who own capital receive in-

come, save part of it, and own even more capital. Section A5 of the appendix shows how these accounts fit together.

Spending and Saving

Think of a bar near a port. If the dockworkers earn less, they spend less at the bar, and the bar takes in less. Those who receive the profits of the automated port may spend elsewhere, and spend a different share of their income.

The share of income spent on consumption is called the propensity to consume. Its simplest form is:

$$c = \frac{C}{Yd}$$

C is consumption, Yd disposable income. This c is an average ratio, different from how much of an extra euro is spent, which is called the marginal propensity. In the model, groups have different average coefficients, and the choice matters for total demand.

A household that spends ninety out of a hundred has a propensity of 0.9. Another that spends sixty and saves forty has 0.6. Shifting ten euros of income from the first to the second can reduce total consumption, all else being equal. The English economist Joan Robinson summed up the thinking of the Polish economist Michal Kalecki in a phrase that became famous: workers spend what they get, capitalists get what they spend.

But what happens to saving? It can finance investment, sit in financial assets, or find few productive opportunities. Saving does not vanish from the economy, but neither does it automatically become new demand able to offset every fall in consumption.

There is also a word that changes meaning. In the previous chapter demand referred to labor; here it refers to goods and services. A firm can demand less labor per product while its customers demand more products. Or it can cut labor and find itself with

fewer customers. In the national accounts, demand for goods and services is written like this:

$$AD = C + I + G + NX$$

C is consumption, I investment, G government purchases, NX net exports. Transfers to households, such as pensions and benefits, are not government purchases of goods and services: they enter household income, which households then spend on consumption. They must not be counted twice. In the model, demand includes consumption, investment and government purchases, and the link between regions runs mainly through capital income that moves; world trade is not represented. Chapter 10 explains the model's version.

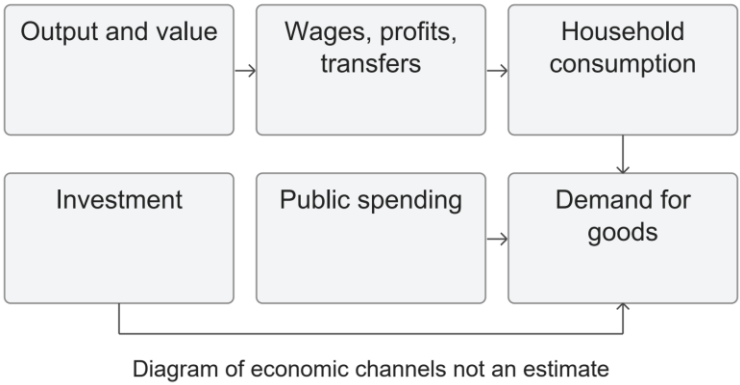


Figure 9 – Income paths within the model framework. The diagram includes output, incomes, consumption and investment, and the role of transfers.

Potential output and output actually sold can diverge. In the model, investment responds to the profit share; consumption, transfers and government purchases play their part in demand. Production generates income, income generates consumption, and the calculation solves all these steps together. If demand falls

short, output drops below capacity relative to the world without AI. Section A6 of the appendix makes the link readable.

Why Demand Does Not Collapse

At the start of this research there was a chain that seemed logical: automation reduces labor income, lower incomes reduce consumption, lower consumption puts profits in crisis. It is the modern version of the Ford story. Its strength is measured against the other channels: new investment, cost savings, cheaper services, public spending.

And the model's result is clear: insufficient demand can reduce output relative to capacity, but in the central results it produces no collapse. What emerges, more clearly than any fall in output, is something else: the concentration of benefits and the difficulties of those looking for work. The sensitivity analysis in chapter 16 confirms it: the parameters for investment and demand transmission explain a small part of the variability in the results.

There is also a technical choice to declare. The version without AI follows a path close to potential, and the demand difficulties of the version with AI are measured relative to that path. It is a closure choice for the model, explained in the next chapter.

To see who wins and who loses, the model divides the population of each region into three groups: the half with the lowest incomes, the middle forty percent, the richest ten percent. The shares of income and wealth are used to build representative groups. Within the poorer half there are people with very different jobs and protections, and the average size of an effect does not say how much it weighs for someone who loses their job. That is why the chapters on work and those on distribution must be read together.

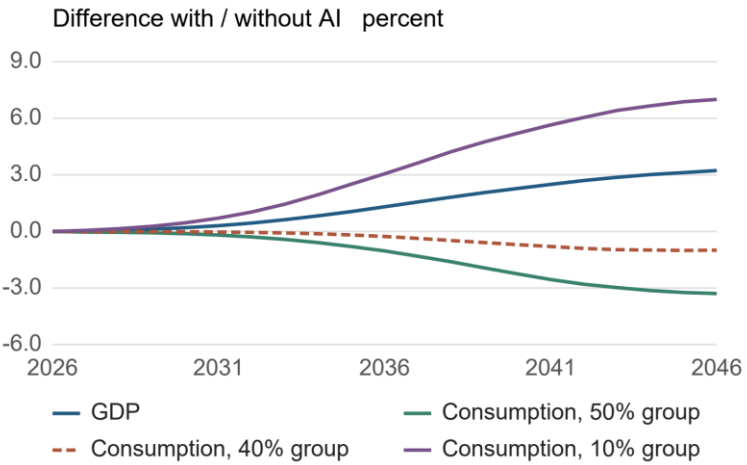


Figure 10 – Original macroeconomic model, medians of the paired differences relative to the world without AI. GDP measures output; the three group curves measure the change in modeled real consumption per capita, used in the model as an indicator of purchasing power. Percentages relative to the benchmark, 2026–2046. Model simulations.

In 2036, world GDP in the median simulation is about 1.3 percent higher than in the world without AI. Consumption of the poorer half is about 1 percent lower. In 2046 GDP is 3.2 percent higher, and consumption of the poorer half 3.3 percent lower.

Be careful how you read this. It does not mean that in 2046 that half consumes less than in 2026. The percentages compare two worlds in the same year. You can be better off than twenty years earlier and worse off than you would be without the new wave of AI. It is the same distinction used for labor demand. And consumption is not all of well-being: free time, job quality, access to services and economic security count too, and chapter 14 accounts for them separately.

Benefits can also reach households through lower prices, not only through higher incomes. It is an important channel, and it

must be measured: it is not enough to list the services AI might make cheaper and give them an arbitrary value.

And where income arrives matters. A firm can use a technology in one country and pay for services or owners in another. In the model, a parameter measures how much of the capital income created by AI is concentrated in certain regions. It does not reconstruct the tax bases of multinationals, but it explains something important: why a territory can go through a labor transition without receiving a corresponding share of the new income.

Larger output can come with losses for some groups. That is why the book treats distribution as one of the mechanisms that drive consumption, work and public finances, not as a detail to add after growth. In the policy chapters you will see which interventions change it and what they cost, always distinguishing the protection of an income from the creation of a job.

First, though, we need to go inside the machine that produces the numbers. The next chapter shows how two comparable worlds are built, how parameters are chosen, and why four thousand simulations are not four thousand observations of the future.

Chapter 6

The Model and the Paired Simulations

In 1747, aboard HMS Salisbury, the Scottish physician James Lind took twelve sailors sick with scurvy, divided them into six pairs in conditions as similar as possible and gave each pair a different remedy: cider, vinegar, sea water, oranges and lemons, and others. After six days the two sailors who had eaten citrus fruit had almost recovered. It is considered one of the first controlled experiments in history. The idea is simple and very powerful: to know what effect something has, you must compare it with something identical in every respect except that one thing.

This book's model does the same with artificial intelligence. But there is an extra difficulty: the term of comparison cannot be a world standing still. Comparing an economy with AI to a motionless past would attribute to the technology the ageing of the population and the progress that would have happened anyway. That is why the book builds two paths that start from the same conditions and change over time, and makes the only difference explicit: the new wave of AI.

What the Model Contains

A model chooses which parts of reality to represent in order to answer a question. This one follows firms, households and the state, in three regions and with three income groups. It moves forward year by year, from 2026 to 2046: technology changes tasks and productivity, and these changes reach wages, employment, consumption, investment and public accounts.

Table 12 – Components of the model.

Component	Scope
Regions	Advanced, emerging and low-income economies
Households	Groups of 50%, 40%, 10% in each region
Production	Technology, labor demand and prices
Government	Taxes, transfers and debt
Time	Annual steps from 2026 to 2046

The model's categories bring together people with very different experiences. The advantage is that the links in the economy can be followed; the price is that an income group or a region does not tell you how each person is doing. Chapters 2, 4, 5, 10 and 11 follow the mechanisms, with references to the full relationships in the appendix; chapter 12 explains how to redo the calculations. Here we look at how the two worlds are built and how their difference is read.

The first thing to know is that in the model machines do not all arrive on the first day. Adoption grows along a curve that depends on three things: the maximum level reachable, the year in which diffusion is fastest, and the speed. Technically it is a logistic curve that starts from zero in 2026.

Technical capability comes before use. An exposed task can be replaced or only assisted; new capabilities and robotics widen exposure over time. The saving on automated tasks raises productivity in proportion to their weight in production: a large improvement in a small activity does not become the same improvement in the whole economy. New human tasks arrive with a delay, and how many there are relative to those replaced, and after how long, decides much of the results.

Then you need to understand how many people are really needed. If machines do a larger share of activities, producing the same quantity requires less human work. If output grows, part of the work comes back. The net effect depends on both. In the model, labor demand is built from output and the human tasks that remain. The wage of the world without AI serves as the reference, and actual wages change the requirement according to an elasticity: how much the labor demanded changes when its cost changes.

An earlier version derived labor by dividing the wage bill by the wage, and so closed the gap between demand and supply too easily: by construction, work always came back. Chapter 15 tells the story of this correction.

The Starting Conditions

The initial conditions mix observed data, adaptations of data and assumptions. Calibration assigns values to the model: some come from published statistics, others approximate a region, others represent possibilities that are still uncertain. The table shows some of the ranges.

Table 13 – Parameters and ranges.

Parameter	Range	Interpretation
New-task lag	4–10 years	Delayed compensation
New tasks ψ	0.25–0.85	Per replaced unit
Wage elasticity ε	0.15–0.75	Demand relative to the wage
Wage response ϰ	0.05–0.50	To the unemployment gap
Automation θ	0.25–0.70	Share of exposed tasks that can be replaced
Central adoption year	2030–2040	Technological diffusion

New tasks arrive after a period of between 4 and 10 years. For every task replaced, between 0.25 and 0.85 new ones arise. Labor demand responds to the wage with an elasticity between 0.15 and 0.75. Wages respond to unemployment with a coefficient between 0.05 and 0.50. The share of exposed tasks that is automated ranges from 25 to 70 percent, and the central year of diffusion falls between 2030 and 2040. Chapter 12 presents them all, with their sources.

Each uncertain parameter is drawn at random within its range, with the same probability for every value, and independently of the others. The bands in the charts therefore show how much the results change within these ranges; demography and some initial values stay fixed.

Twin Worlds

Here is the heart of the procedure. For each set of conditions, the world without the new wave of AI is calculated first, then the

world with AI, with all the other parameters identical. The results for the same year are subtracted within the pair. Only afterwards are the differences from all the pairs put together. Like Lind's pairs of sailors: first you compare within the pair, then you look at what happened in all of them.

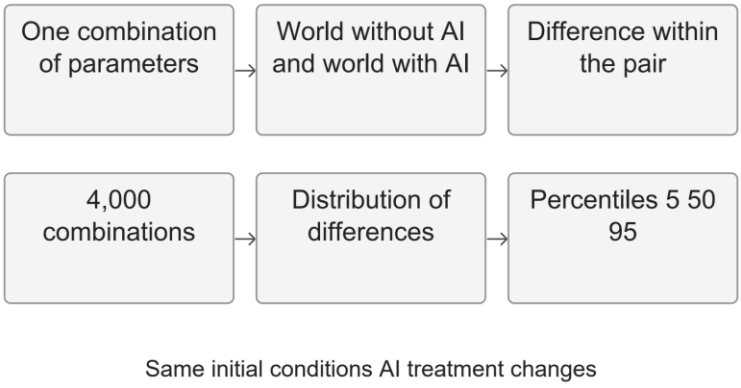


Figure 11 – Procedure of the paired simulations. First the difference within each pair is calculated; then the set of differences is summarized. The world without AI keeps ordinary technical progress and the other conditions set out in the model.

For the book's charts the procedure was repeated 4,000 times, with the numerical seed 20261010. Code, configuration and seed allow anyone to regenerate exactly the same draws. The pair removes arbitrary differences between the two worlds; an error present in both, however, also remains in the difference.

Four thousand pairs give four thousand differences, and they must be summarized. The median is the value in the middle when you line them up; the mean is obtained by adding them up and dividing by the number of values. And a simulation in the middle for GDP may not be in the middle for unemployment: the medians of different indicators do not necessarily describe the same world. Percentiles indicate other positions: the fifth leaves about 5

percent of the values below it, the ninety-fifth about 95. Between the two lies the central 90 percent of the draws for that indicator.

Table 14 – Differences in 2036 relative to the world without AI, original macroeconomic model. GDP and consumption percentages are relative changes; unemployment is expressed in percentage points. Percentiles of 4,000 draws, not empirical confidence intervals. Model simulations.

Difference in 2036	P5	Median	P95
GDP %	0.3	1.3	3.5
Consumption, 50% group %	−4.9	−1.0	0.1
ADV unemployment pp	1.4	4.3	10.7

In 2036 GDP is 1.3 percent higher in the median, and 90 percent of futures lie between 0.3 and 3.5. Consumption of the poorer half is 1 percent lower in the median, and ranges from minus 4.9 to plus 0.1: in the best future in twenty, the poorer half gains barely a tenth of a point. Unemployment in the advanced economies is 4.3 points higher in the median, and ranges from 1.4 to 10.7. Even in the most favorable future in twenty, it rises. Chapter 14 goes into more detail.

You can also count how many draws meet a criterion: that is a frequency. Chapter 21 uses frequencies to give futures a name: inclusive, polarized, discontinuity. The age module, instead, studies some specific cases, such as slow retraining: what happens if work comes back slowly and finding a job is hard? The answer serves to choose which data to collect and which policies to examine.

Accounts That Balance

The model follows payments, savings and investment among households, firms, states and regions, and every euro must have a sender and a recipient. In the reproduction made for the book, the largest discrepancy in the accounting identity is of the order of 10

to the minus 16: the accounts close to within the computer's rounding. The checks of the age module also verify that people are conserved from one year to the next, that flows add up and that no group becomes negative.

Accounts that balance say the calculation is consistent; behavior is checked against history, external evidence and results not used in the calibration. Chapter 18 says which.

To understand which assumptions move the results most, the book uses Sobol analysis, which separates the contribution of each parameter and of their interactions. What matters most is the pace and breadth of automation, the return of human tasks and the way the labor market adjusts. Chapter 16 reads it calmly.

The age module receives demand from the macroeconomic model and follows the employed, job seekers and the discouraged. Chapter 26 closes the circle: it feeds unemployment back into wages and benefits, and recalculates everything.

The Assumptions, Once and for All

Here are all the model's assumptions, in one place. They are not repeated in the rest of the book: every result holds within this list.

The structure. Three regions: advanced economies, emerging and middle-income economies, low-income countries. In each, three income groups: the poorer half, the middle forty percent, the richest ten percent. One year at a time, from 2026 to 2046. No inflation, exchange rates or trade between countries; the interest rate is given. The regions are linked by AI capital income: between 20 and 60 percent of that arising in emerging and poor economies is booked in the advanced ones.

The technology. Starting exposure is the one estimated by the International Monetary Fund: 60 percent of employment in the advanced economies, 40 in the emerging ones, 26 in poor coun-

tries. It grows by up to 60 percent more by 2046, without exceeding 95 percent. Of exposed work, between 40 and 70 percent is substitutable; of the substitutable part, between 25 and 70 percent is automated by 2046. The model does not contain an AI able to do almost everything.

Adoption. An S-curve starting from zero in 2026, with the year of fastest diffusion between 2030 and 2040 in the advanced economies and a slope between 0.25 and 0.60. Emerging economies follow 1–5 years later, poor countries 3–10 years later.

Productivity. The saving on automated tasks is worth between 15 and 45 percent of their cost; the help on assisted tasks between 0 and 30 percent. New tasks enter growth with a weight of one half. Without AI, the labor content of tasks erodes anyway by between 0.1 and 0.35 percent a year.

Work. For every task replaced, between 0.25 and 0.85 new ones arise, after 4–10 years. Labor demand comes from output and human tasks, with a wage elasticity between 0.15 and 0.75. Human and automated tasks substitute for each other with an elasticity of one. Wages follow productivity with a pass-through between 0.3 and 1 and respond to unemployment with a coefficient between 0.05 and 0.50. Bargaining returns to workers between 0 and 30 percent of the share lost. Labor supply follows United Nations demographic projections.

Demand. The world without AI runs at potential. With AI, a demand shortfall passes to output with a coefficient between 0.3 and 1, and investment responds to the profit share with an elasticity between 0 and 1.

Distribution and the state. Income and wealth shares from the World Inequality Report; propensities to consume that differ by group; hypothetical fiscal rules, the same in both worlds; interest rates on debt that can move up or down by one point.

The age module. Four age groups, 15–24, 25–34, 35–49 and 50–64. Probabilities of first entry, return and discouragement and accessibility weights by group set by the book; the advanced economies start from the situation of the European Union.

Observed data. The labor series, Stanford/ADP for the United States and Eurostat for Italy and Europe, are read with a weight for each year: 100 percent for 2026 and 2025, 50 for 2024, 5 for 2023, 1 for 2022. 2022 and 2023 reflect post-pandemic hiring and higher interest rates. The complete series remain in the book.

Sampling. Twenty-one parameters drawn uniformly and independently within their ranges, four thousand pairs of worlds, seed 20261010.

What is left out. Differences within income groups; ages in the main model, which the module adds separately; energy costs; benefits that consumption does not measure, treated separately with a parameter between 0 and 5 percent of consumption.

The forecast by occupation in chapter 36 goes beyond this model and has its own rules, listed there once. First, though, the next chapters look at people: who loses their job, who gets in, and how long the wait lasts.

Chapter 7

Outflows and Re-employment

In the 1980s three economists, Louis Jacobson, Robert LaLonde and Daniel Sullivan, followed thousands of workers in Pennsylvania who had lost their jobs in mass layoffs after many years with the same company. Many found work again, some in firms similar to the ones they had left. But for those who left a firm in distress after so many years, the long-term loss averaged about a quarter of their earnings every year, compared with similar colleagues who had kept their jobs. The study, published in 1993, became a classic for a simple reason: it showed that finding a job again does not erase the cost of having lost one.

A month and three years without income are different experiences. And even a quick return can bring pay, hours or stability different from before. This chapter looks at how the model represents exits from work, hires and the time needed to return to work.

Work Photographed and Work Lived Through

The number of people employed on a given date does not say which paths led there. Two economies can have the same total and opposite stories: in one almost everyone keeps their job, in the other many lose it and others find one. For incomes, families and skills, these are different worlds.

To the concepts of stock and flow seen in the previous chapters, time must be added. Someone who loses a job and returns to work after a month does not live the same experience as someone who returns after three years. And not even a quick return guarantees that wages, hours and stability are unchanged. To describe a transition you need a starting situation, the moves from one state to another, and the characteristics of the destination.

The macroeconomic model works mainly with aggregate annual quantities. The age extension, instead, records ordinary exits, cuts

due to a lack of jobs, hires and exits from active search. Individual biographies and contracts require other data.

Table 15 – Distinct movements. An exit from employment does not by itself identify an involuntary loss; exits from the model's scope require a separate reading.

Movement	Meaning in the module
Ordinary separation S	Assumed turnover with no individual reason
Additional cut K	Remaining employed beyond job capacity
Exit at 65	Assumed boundary of the 15–64 scope
Participation exit X	Employed beyond updated capacity

A termination recorded in administrative archives is something else again. It may concern a contract that ends while the person holds a second one, or a change of employer without a single day of unemployment. Counting terminated contracts and calling them "people pushed out of work" would make the phenomenon look larger or smaller depending on how the perimeter is drawn.

How a Loss Arises in the Model

In the age module, each year follows a precise order. First people age by one year. Then the number who can take part in work is updated. Then some of the employed leave through normal turnover. At this point the program compares the remaining employed with the labor demand assigned to their age. If they exceed the jobs, the difference becomes a cut. Only afterwards does the matching between remaining vacancies and job seekers take place.

The cut is written like this:

$$K = max(0, E\,before - D)$$

E before is the employment remaining after the previous steps, D the accessible jobs in the same group. The max function picks the

larger of zero and the difference. If ninety employed remain for eighty jobs, the cut is ten. If seventy remain for eighty jobs, the cut is zero, and the problem becomes who will manage to fill the ten vacancies.

The order of operations has important consequences. If ordinary exits already reduce employment enough, weaker demand shows up as fewer hires, not as additional cuts. The program does not turn every fall in demand into a dismissal. And the same happens in reality: a firm that needs to reduce staff often does not fire anyone. It waits for someone to retire or leave, and does not replace them.

That is why a series of cuts close to zero can coexist with lower employment than in the world without AI: the deterioration passes through unreplaced jobs, harder entry, slow re-employment.

Before looking at exits, it is worth seeing the overall result. The figure shows the employment rate of the four age groups in the advanced economies, in the scenario where retraining is slow. Each panel compares the version with AI and the version without, with the same baseline conditions within each pair.

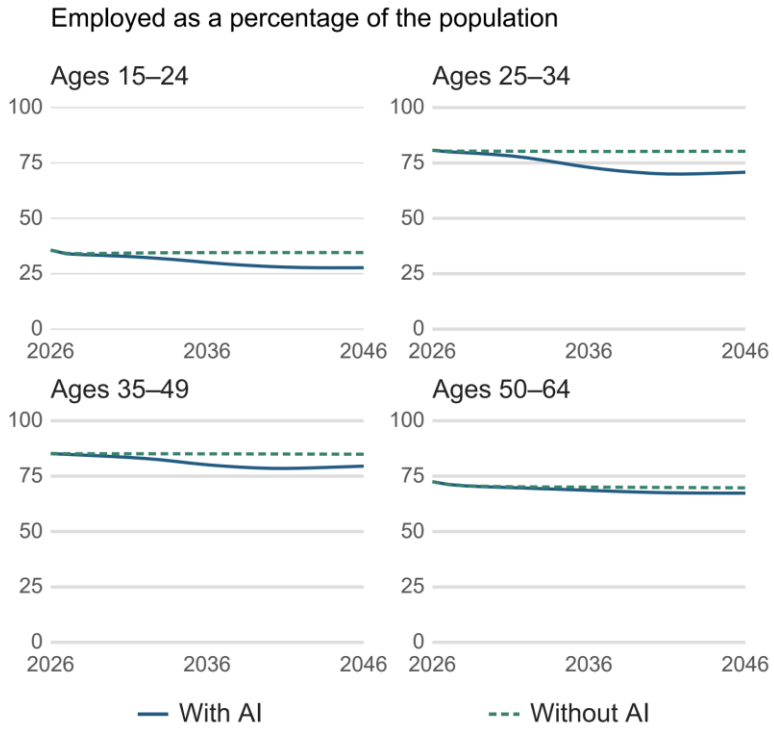


Figure 12 – Simulated employment rate, advanced economies, slow-retraining scenario. Percentage of the population of each age group. Medians of 4,000 draws per version. Vertical scale common to the panels. The age groups describe ages in each year, not four unchanging groups of people. Initial EU profile used as a proxy; AI and matching parameters assumed.

The employment rate relates the employed to the whole population, so it shows a difficulty even when some people stop looking. The differences between ages depend on the initial conditions, the accessibility weights and the entry probabilities; exposure to AI depends on occupations and tasks, not on age as such.

Someone who is twenty-four in 2026 is thirty-four in 2036. The model ages people, and the 15–24 panel gathers, each year, who-

ever is that age at that moment: it follows an age group, not the same people.

In 2036 the employment rate of 25- to 34-year-olds is about 6.8 points lower than in the world without AI, as the median of the differences within the pairs, calculated as chapter 6 explains.

Few Layoffs, Fewer Hires

Employment can also fall through exits that are not replaced. In that case the decline shows up in hiring before it shows up in a wave of cuts. The next figure compares two kinds of simulated exits between ages 50 and 64. Ordinary separations are normal turnover: early retirements, resignations, contracts that end. Cuts, instead, are the additional reduction imposed when the remaining employed exceed the available jobs.

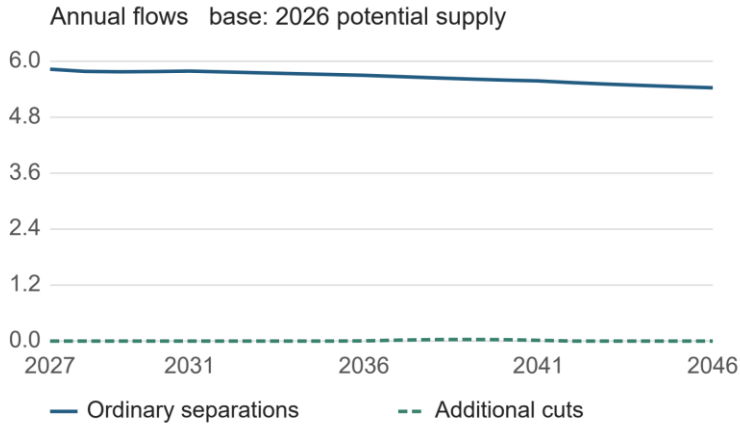


Figure 13 – Simulated ordinary separations and additional cuts, age group 50–64, advanced economies, slow-retraining scenario with AI. Annual flows as a percentage of the age group's potential supply in 2026. The first year shown is 2027: 2026 initializes the states and does not measure flows. Medians; the cuts curve is almost zero on the scale of ordinary turnover.

In the results file, the median of cuts in the first three age groups is zero in every year of the slow scenario. Between 50 and 64 the maximum annual median is about 0.035 percent of initial potential supply. In some simulations there are cuts, but in the median the fall in employment passes almost entirely elsewhere.

The reason is turnover. Every year some of the employed leave their jobs, and demand can shrink by absorbing these exits. If many leave and few are replaced, employment falls without a large wave of cuts. It is the quietest form of a labor crisis: nobody is sent away, simply nobody is called in anymore. Attention therefore goes to hiring.

A small median cut does not mean the problem does not exist: it means the problem takes the form of missing hires, not layoffs.

There is also an effect that looks like a paradox. If ordinary separations are a share of employment, an economy with fewer employed people also has fewer separations in absolute terms. That does not mean it protects workers better: the base on which the share is calculated has simply shrunk.

The same problem arises when comparing territories of different sizes. A flow measured on employment at the start of the year answers one question; a flow measured on 2026 potential supply, as in the chart, answers another. The book uses the second base to keep all movements comparable; the risk faced by those already employed is read with the first.

Accounting for People

So that an exit due to age does not look like an expulsion from work, each year's balance must use the same perimeter. With ages already updated:

$$E\,final = E\,updated - X - S - K + H$$

E updated is employment after the ageing step. X collects the exits due to participation limits, S the ordinary separations, K the cuts, H the hires. E final is what remains after all these movements.

The exit at sixty-five happens earlier, when people age. It is the boundary of the 15–64 perimeter chosen by the module, not a retirement age for everyone.

Good accounting prevents people from vanishing inside formulas. But it does not say what they have lost. For a family, what may matter is the income between one job and the next, a forced move, the cost of a course, turning down a contract incompatible with caring for children or parents. The book measures them when it comes to policy.

The Pennsylvania study is a reminder. Returning to work closes a period without work, but it does not restore the conditions of before. To understand a return you need the length of the wait, the wage, the hours, the stability of the new contract. And to link the job lost and the job found you need data that follow the same people over time.

The next chapter uses hires and re-employment probabilities to look inside this gap. And it meets those who have never had a first job: people trying to start. For them the question is not how to recover a lost job, but how to gain the experience needed to get one.

Chapter 8

Entering Work for the First Time

"Wanted: young person with experience." It is the oldest joke in the labor market, and nobody aged twenty finds it funny. Experience is often the requirement for being hired. But it is built precisely through the first job.

For decades the circle was broken in a simple way: firms hired beginners for easy tasks and trained them there. The young clerk at the law firm who photocopied and filed, the trainee who prepared the first draft, the new call center hire who answered the most basic questions. If artificial intelligence does the easy tasks, that opening narrows. This chapter tries to measure by how much.

And it is no small matter, because a difficult entry does not cost just one year. Research on young people who enter the labor market during a recession shows this clearly. Lisa Kahn followed American college graduates who left university between 1979 and 1989: those who graduated when the economy was worse started on lower wages, and the difference remained visible throughout the period observed, many years later. Oreopoulos, von Wachter and Heisz, studying Canadian graduates, found that a severe recession, with unemployment five points higher, initially cost about 9 percent of earnings; the loss halved within five years and disappeared after about ten. Ten years. The moment you enter leaves a mark.

Two Different Waits

A young person looking for a first job and a forty-year-old looking for a new one appear in the same unemployment statistic. But they are not in the same situation. The second has already been through at least one job: they have a CV, contacts, a trade a recruiter recognizes. The first still has to get their first chance. Their obstacles are different: skills, networks, signals that employers can read, willingness to accept certain conditions.

And age is not enough to tell them apart. A twenty-three-year-old may already have changed jobs three times; a thirty-five-year-old may be looking for the first one after long studies or other activities. That is why the book does not call every hire of a young person a "first entry" and every hire of an adult "re-employment". What matters is where the person comes from, not their date of birth.

The module that follows ages divides job seekers into two groups. Group F gathers those looking for their first job. Group U those looking after a previous job. From each group, every year, comes a flow of hires: first-entry hires and return hires. Keeping them separate reveals different movements that total hiring would hide.

$$H = H\,first + H\,return$$

H is the total of simulated entries into work in the year. It does not include those who move directly from one employer to another without becoming unemployed.

The young people who reach working age each year are divided between the two groups with fixed shares: "first entry" is a category of the model, which comes close to the first contract of a person's life without coinciding with it.

How wide is the door? The number of hires depends on how many people are waiting and how easy it is to get in. A flow can grow because the queue grows, even if each person has a lower chance of being hired. Or it can fall because the queue gets shorter, while those who remain get in more easily. To understand the door you need both the flow and the right denominator. The first-entry rate is measured like this:

$$First - entry\,rate = \frac{H\,first}{Eligible\,F} \times 100$$

The denominator is the young people in group F looking for work before they meet the available jobs in the year. The numerator is how many of them are hired. Students who are not looking and those not in the queue are not in the denominator.

In the code, hires are built in two steps. First an offer of hires is calculated, starting from baseline probabilities, a factor for matching people and jobs, and a barrier that grows with AI adoption. Then the jobs actually available set a ceiling. A high probability does not produce unlimited hires when jobs are scarce. And those seeking a first job and those returning compete for the same jobs.

Table 16 – Teaching example. Hires increase from twenty to thirty, but the rate among job seekers falls from forty to thirty percent because the pool doubles. They are not model data.

Measure	Before	After
Eligible job seekers	50	100
Hires	20	30
Entry rate	40%	30%

In the example, hires go from twenty to thirty: it looks like good news. But the young people in the queue have doubled, from fifty to a hundred, and the probability of getting in falls from 40 to 30 percent. A city can announce more hires and have a longer queue outside the door. A group can obtain fewer contracts simply because it is smaller, without selection having become tougher. Whenever you look at a chart about work, ask yourself whether you are looking at a count, a flow or a rate among those actually searching.

What the Model Says

The next figure shows first-entry hires of young people aged 15 to 24 in the advanced economies, in the scenario where retraining is

slow. The two curves use the same base: all young people in that age group available to work in 2026. That way volumes are compared without passing off indices as real people.

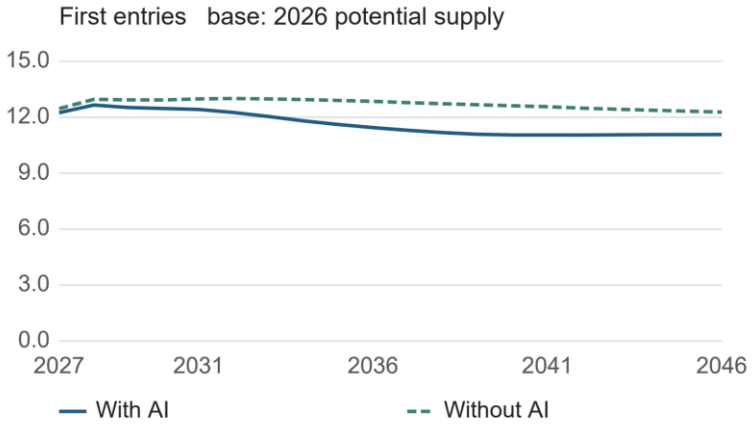


Figure 14 – Hires in the first-entry category, age group 15–24, advanced economies, slow retraining. Annual flows as a percentage of the age group's potential supply in 2026, from 2027 to 2046. Medians of the paired versions with and without AI. The classification of initial and new job seekers is assumed; these are not observed first contracts.

In 2036 the difference between the world with AI and the world without it is about 1.37 fewer hires for every hundred young people in the initial base. Put like that, it seems small. But look at it from the point of view of those in the queue. In the same year, the probability that a young person seeking a first job is hired during the year is about 14.3 percentage points lower. The two numbers do not contradict each other: they have different denominators and tell two sides of the same story. The first measures how many entries are missing relative to all young people. The second measures how much the door has narrowed for those standing in front of it.

Fourteen points, for someone looking for a first job, are the difference between a door left ajar and a door that is closing.

The scenario has new tasks that arrive late, human labor that is not fully recreated, higher barriers for those without experience. The world without AI has the same starting conditions: the difference between the two is the effect of AI on entry. It is an average for the advanced economies, not the number for every country or every occupation.

Chapter 3 showed the American figure: 22- to 25-year-olds in the occupations most exposed to AI fell by 4.4 percent in the last year, more than twice as much as their peers in less exposed jobs, through hiring that did not happen. It is a different measure, employment in a payroll sample. But the model and reality, starting from different roads, point to the same place: the entrance door.

There is also something the module does not see, and it makes the problem more serious than it calculates. A first job does not only produce a result: it produces experience. The beginner learns to recognize an anomaly, to understand a customer, to coordinate with others. If those tasks are automated, the organization has to decide how it will build the experience needed for everything else. Technology can help, with exercises, simulations, assistance. But having the tool and learning on the job remain two different things.

The current module does not include this accumulation of skills. In the model, a lower probability of entry does not worsen the future abilities of those who wait: there is no variable that turns waiting into lost experience. Studies of those who enter during a recession say that this loss exists and lasts for years. The real effects, therefore, last longer than the model calculates.

The Assistant Paradox

Let us return to the call center of the first chapter. Brynjolfsson, Li and Raymond, in the 2023 version of their study (Generative AI at Work), followed the gradual introduction of a conversational assistant among 5,179 customer support agents. Problems resolved per hour rose by 14 percent on average, and by 34 percent for beginners and the less skilled; for the most experienced the effect was small. Source: https://doi.org/10.3386/w31161.

It is good news for those already inside. But here there is a paradox. The more the assistant makes beginners good, the fewer beginners are needed to do the same work. And the more it makes experts autonomous, the less need there is to pair them with young people. The same tool that helps those who got in can make it harder for those outside to get in.

The study could not measure this: it looked at people already hired. Young people who were never hired do not appear in any company record. They are the invisible part of every productivity experiment: the benefit for those already working and the difficulty of those who do not get in are measured separately.

And resolving more problems per hour raises a question about the contract: do pay, hours and pressure change? It is the firm's decisions and the sharing of the gain that link technical capability to the conditions of those who work.

Finally there is the question the study leaves open, and the one that matters most for a young person. Does a beginner who works well with the assistant really learn? Could they do the same job if the assistant were switched off? Would they recognize a wrong answer? Would they learn things they could take to another job, or just get used to a program? Chapter 24 shows an experiment in a school that gives a worrying answer.

The same applies to recruiters. A result depends on the person, the tools and the supervision. Excluding those who use AI would be absurd. Automatically attributing to the person all the capabilities of the tool would be equally wrong.

The Door for Those Returning

Those who have worked before go through another door: the door of return. Here the denominator is the job seekers in group U, before they meet the year's jobs.

$$Return\ rate\ = \frac{H\ return}{Eligible\ U} \times 100$$

An improving rate can coexist with many people still outside, if the queue grows. A worsening rate can remain invisible in total employment for a while, because those with a job keep working. That is why the book shows re-employment separately from employment and from first entry.

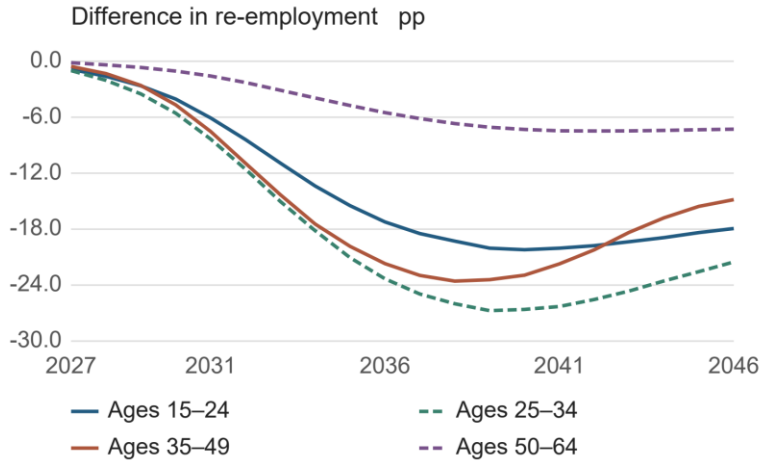


Figure 15 – Difference with AI minus without AI in the annual re-employment rate, advanced economies, slow retraining. Percentage points of eligible job seekers in category U; medians of the differences in the 4,000 pairs, from 2027 to 2046.

In the slow scenario, in 2036, 25- to 34-year-olds have a probability of re-employment about 23.3 percentage points lower than in the world without AI. Not because they have become less capable or less willing. The result depends on the matching of people and jobs and on how many jobs there are. When a system generates few accessible jobs, blaming those who "fail to update their skills" turns a problem of the economy into a judgment on the person.

And even those who return do not tell the whole story. The code records that a person goes back to work. It does not know at what pay, how far from home, for how many hours, or whether the new job uses their abilities. Re-employment can improve a precarious situation or replace a good lost job with a worse one. The rate answers the question "how many return?", not "how do they return?".

There is also a trap in averages. Those with the skills most in demand find work first. Those still in the queue the following year are, on average, those who struggle most. So the rate can worsen even if the door has not changed: the queue has. The module simplifies all this into age and probabilities; in real data, separating the barrier from selection requires experience, education, the occupation sought and the territory.

Finally, the module distinguishes the jobs that exist from those that actually get filled. There can be vacancies and people queuing at the same time, if the match fails.

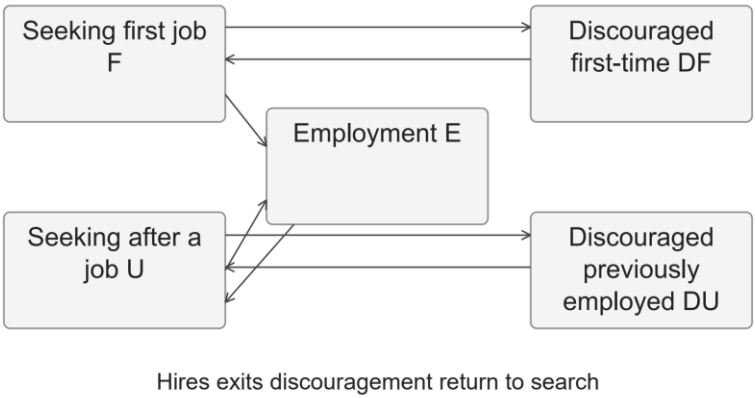


Figure 16 – States and movements represented by the module. Two categories of job seekers can enter employment or discouragement; from the latter they can return to search.

The diagram summarizes the paths: those seeking a first job, those searching after having had one, those working, those who become discouraged and those who start searching again.

Filling the diagram with real data requires long work histories: the Eurostat transitions series distinguishes employed, unemployed and inactive, but has no ages.

One thing, however, needs no diagram. An entry that comes late delays everything else: experience, pay, independence, a home, children. Studies of recessions say the mark lasts for years.

The next chapter tackles the answer everyone gives to this problem: retraining. Just train people for the new jobs. But when do the new jobs arrive? And how do those who wait live in the meantime?

Chapter 9

Occupational Retraining

Since 1962 the United States has had a program designed precisely for people who lose their jobs because of a major economic change: Trade Adjustment Assistance, created for workers hit by import competition. It offers training, income support during courses, and help looking for work elsewhere. In 2012 a national evaluation of the program, carried out by Mathematica Policy Research, compared participants with similar workers who had not taken part. Over the following four years, participants had earned less on average than their peers. Part of the explanation is simple: while they were in training they were not working, and when they finished they had to start from scratch in a new trade.

This does not mean training is useless. It means that retraining requires two timetables to meet: the time needed to learn, and the time within which a suitable job becomes available. A course can end before the new jobs appear, or arrive when lost income has already made it impossible to attend. To evaluate a transition you have to follow both sequences.

What Retraining Means

The word retraining covers different things. A firm changing its business. An occupation changing its tasks. A person learning something new. A territory looking for other ways to produce. These movements are connected, but they do not necessarily happen together. A course can end before a suitable job exists; a job can exist before those looking for work can reach it.

The model distinguishes two mechanisms. The first is the creation of new human tasks, which offset part of those displaced: it changes how much labor is needed, in the macroeconomic model. The second is the matching between job seekers and accessible jobs: it decides how much of that demand actually becomes em-

ployment, in the age module. Confusing them would make automatic a transition that the equations keep separate.

The first mechanism has two numbers: how much is recovered and with what delay. The parameter ψ says how many new tasks appear for each task displaced, in the model's relationship. The delay says when recovery begins to be felt.

The second mechanism uses baseline probabilities, job limits and factors linked to AI adoption. It contains no catalogue of courses, destination occupations or moves from one city to another.

The module compares four cases. The reference scenario keeps the model's original ranges and adds flows by age. Effective compensation adds more new tasks and easier matching. Slow retraining reduces recovery and lengthens delays. Persistent displacement combines weak recovery, broader automation and harder matching.

Table 17 – Assumptions configured in the project, not empirical estimates. ψ describes new tasks per unit of substitution; the lag is in years; the matching factor multiplies the baseline probabilities. The scenarios also change other conditions stated in the code.

Scenario	ψ	Lag, years	Matching
Reference	0.25–0.85	4–10	1.00
Effective compensation	0.85–1.20	2–6	1.20
Slow retraining	0.25–0.65	6–12	0.80
Persistent substitution	0–0.35	8–16	0.65

In the effective scenario ψ can exceed one: the model allows the tasks recovered to outnumber those displaced. Jobs, however, do not grow by the same amount: adoption, output, wages and matching also matter. And persistent displacement also changes

the extent of automation and exposure: from one case to another several levers move together.

The Gap That Remains at the End

The figure follows 25- to 34-year-olds in the advanced economies. For each scenario it shows how far their employment rate in the world with AI moves away from that in the world without. Zero means equal rates; negative values, a lower rate with AI.

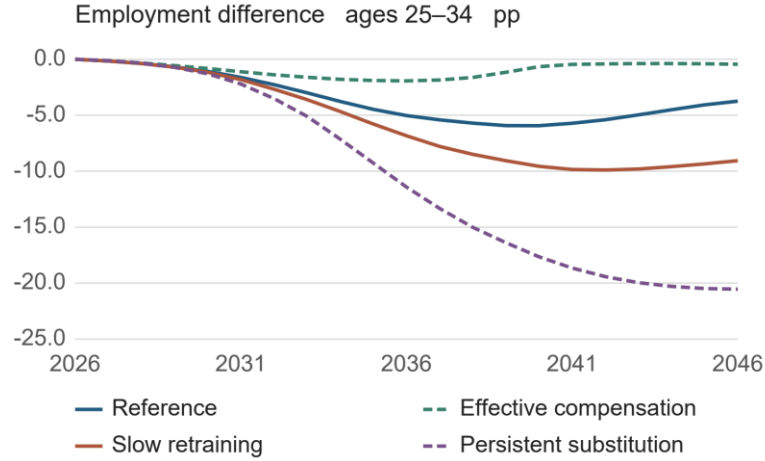


Figure 17 – Difference with AI minus without AI in the employment rate of the 25–34 age group, advanced economies. Percentage points of the age group's population; medians of the paired differences. Four exploratory scenarios, 2026–2046, with no assigned probabilities. Each comparison keeps its own no-AI baseline matching; several parameters change across scenarios.

In 2046 the median difference is about minus 0.4 points in the effective scenario, minus 3.7 in the reference, minus 9.1 in slow retraining and minus 20.5 in persistent displacement. Between the best and the worst case lies the distance between a problem almost absorbed and a generation in which one in five of those who would have worked is left out.

The decisive question is this: does recovery come soon enough, and reach enough people, to offset displacement? The figure shows that almost everything depends on that answer, and therefore where to look for evidence and which policies to prepare. And in 2046, the last simulated year, the difference is still negative in all four cases.

Then there is the question of how long the wait lasts. The re-employment rate of chapter 8 helps to reason about it, but it is not a distribution of durations. An abstract example clarifies the link. If a person has the same probability p of finding work every year, keeps searching and nothing else happens, the probability of still waiting after n years is:

$$Probability\ still\ waiting\ = (1 - p)^{n}$$

With p equal to 0.5, after three years 12.5 percent are still waiting. With p equal to 0.25, 42.2 percent. It is enough for the annual probability to fall from a half to a quarter for the share still waiting after three years to go from one in eight to more than two in five. This is the cumulative effect of a recurring obstacle: a slightly lower probability each year becomes an enormous difference after a few years.

In the module the probability changes over time, people age, leave the search, find limited jobs; the formula serves to see the cumulative effect. And those who stop looking should not be counted as re-employed.

When the Wait Leaves the Statistics

The model keeps those who become discouraged in two groups: those who were looking for a first job and those who had worked before. Active search can fall without employment rising. That is why lower unemployment can coexist with a wider difficulty than

the rate alone shows. To see it, the module defines potential supply and broad non-employment:

$$P* = L + DF + DU$$

$$B = \frac{P* - E}{P*} \times 100$$

L is those actively searching, DF and DU the two groups of discouraged, E the employed. P star is not the whole population: it includes those who take part in work and those who have stopped looking. B measures the part of this set that is not working.

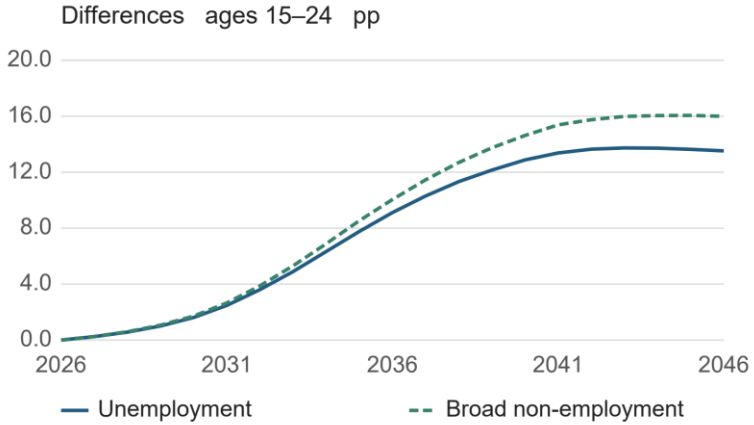


Figure 18 – Age group 15–24, advanced economies, slow retraining. Differences with AI minus without AI in the simulated unemployment rate and in broad non-employment, in percentage points. Different denominators: active supply for the first, potential supply for the second. Medians of the differences within pairs; do not subtract the two curves to count the discouraged.

In 2036, among 15- to 24-year-olds in the slow scenario, simulated unemployment is about 9.1 points higher than in the world without AI; broad non-employment about 10 points higher. Keeping the discouraged in makes visible part of the difficulty that the unemployment rate alone hides.

The module also allows people to move back from discouragement to search. That return lengthens the queue of job seekers, but it is not a hire. A policy that helps people start searching again pushes unemployment up before entries into work increase: viewed on its own, that indicator would mistake an improvement for a deterioration.

Training can be part of the answer. But to evaluate it you have to distinguish enrolment, attendance, learning and entry into a job. Knowing when a course starts does not tell you when someone will be back at work. A completed course is not a job found, and chapter 31 explains how to tell whether a job found is thanks to the course.

And conditions the model does not describe are needed: free time, the ability to move, tools for studying, an income while waiting. Two people equally able to learn can have different fates if one can give up a salary for a while and the other must immediately accept a job that leaves no time for the course. That is what the evaluation of Trade Adjustment Assistance showed: time in training is time without income. These dimensions link work to the distribution of income, the subject of chapter 5.

Table 18 – Different outcomes to observe when assessing a transition. The current module covers only some dimensions; no single row measures the overall success of a policy.

Dimension	Measure to distinguish
Entry	First job and return after a job
Time	Completed waits and periods still open
Quality of work	Pay, hours, stability and match
Participation	Search, discouragement and return
Training	Access, learning and employment

In the policy chapters it will be necessary to say each time which lever is being moved. A transfer can protect income without creating a job. An investment can increase productive capacity without guaranteeing that a job seeker reaches it. A course can improve matching without offsetting labor demand that is too weak. The value of an intervention depends on the problem it addresses and on what is actually measured.

The Job That Stays Empty

Bureaucratic hurdles and the failure to recognize skills can lengthen a transition even when training exists. An electrician who moves to another country may have to retake exams to work. The model compresses these obstacles into parameters.

To understand which of these obstacles weighs most, the book carried out an experiment. It holds the macroeconomic path at its central values and varies four levers of the matching between people and jobs: matching efficiency, from 0.5 to 1.5; the barrier for those seeking a first job, from zero to 0.9; the mismatch between skills and jobs, from zero to 0.8; additional discouragement, from zero to 0.07. The ranges are wide enough to contain all the mod-

ule's cases. The design, known as Saltelli's, starts from 4,096 combinations and requires 24,576 runs of the module, each with and without AI. The question is which lever makes the result vary most.

Table 19 – Sensitivity of the age module: total-order indices, including interactions. Employment: AI minus no-AI differential for 25–34-year-olds; first entry and broad non-employment: 15–24-year-olds. ADV. Central macro path fixed, four assumed ranges. These are numerical estimates of the indices, not shares of people or policy effectiveness.

Assumption	**Empl. 2036**	**Empl. 2046**	**Entry 2036**	**Non-empl. 2036**
Matching	0.6170	0.4088	0.4187	0.2467
Entry barrier	0.4154	0.5219	0.6100	0.7634
Mismatch	0.0282	0.2828	0.0205	0.0245
Discouragement	0.0002	0.0161	0.0004	0.0005

The first-entry barrier weighs most on the two measures for young people in 2036. General matching efficiency weighs most on the employment of 25- to 34-year-olds in the same year. By 2046 the barrier becomes the most important lever for them too, and mismatch gains weight. The message is clear: in the model, what decides the fate of the young is not so much how well the market works in general, but how high the threshold is for those without experience.

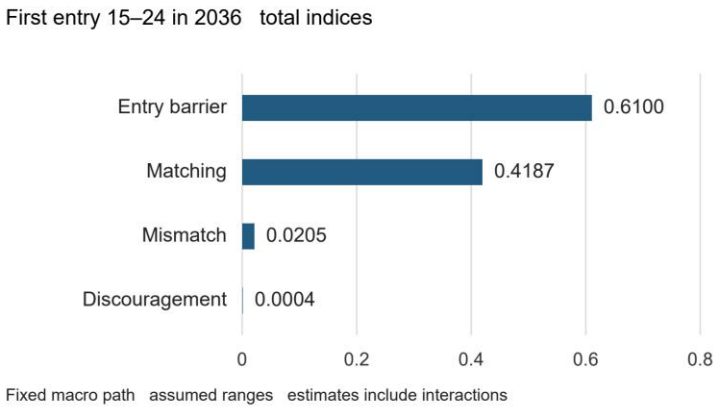


Figure 19 – Total-order indices for the 15–24 first-entry differential in 2036, ADV. Jansen estimates, fixed macro path. The bars include interactions and should not be added up as separate components.

To read these indices, imagine changing one assumption while holding the others fixed, and repeating the comparison across many configurations. The total index compares the variability produced this way with the overall variability, and includes the interactions between levers. A high value indicates an important lever. The formula is in section A7 of the appendix.

And a small index means only that the result changes little when that parameter moves: weighing little in the model's variance does not make a social phenomenon irrelevant.

For contributions close to zero, such as discouragement, numerical precision is limited and the order among them does not matter: for discouragement, on 2036 employment, the first-order index is 0.0018 and the total 0.0002. The large differences between barrier, matching and mismatch, by contrast, are clear-cut.

Those who fear slow retraining and those who believe in effective compensation have the same thing to watch: access times, demand for new occupations, the people who really get into the jobs created.

To study a concrete intervention you change one condition at a time: chapter 24 does this with three policies for young people's entry.

Chapter 26 feeds these trajectories back into the wages and benefits of the main model. Meanwhile one thing is clear: the length of retraining ties the labor market to the resources with which people can afford to wait. The next chapter widens the view to territory: where value is produced, where income arrives, and which services make the transition possible.

Chapter 10

Value Added, Income and Regions

In July 2016, Ireland's statistics office announced that the country's GDP had grown by 26 percent in 2015. Nothing extraordinary had happened in Irish factories, shops or offices. Some large multinationals had moved the ownership of patents and other intangible assets to Ireland, and with them the profits those assets generated. On paper Ireland was producing much more; its inhabitants did not notice. The economist Paul Krugman called it "leprechaun economics", and Ireland had to create a new indicator, modified national income, to measure how much really stayed in the country.

It is the clearest story there is of a truth this chapter follows closely: the value produced in a territory does not necessarily coincide with the income that stays with its inhabitants. Owners and suppliers can be elsewhere; workers can spend somewhere else. To understand what automation does to a place, you have to define the boundary of the account and follow who receives the payments.

Revenue and Value Added

A price says how much is paid in a transaction. But to know how much that transaction adds to output, you have to separate what was bought from others from what the activity added. The value of the goods transported is not the value added of transport: the goods may have been produced elsewhere, while the transport company provides a service. A bar buys coffee, milk, cups, machines, energy. The value of its sales cannot be added in full to every purchase along the supply chain, otherwise some things would be counted several times. Value added serves to avoid that.

Take a service that takes in a hundred and buys forty of goods and services from others. The difference, sixty, is value added. From it labor is paid, the result of the business, and the other

items; the full account also includes taxes, subsidies and the wear of capital. The hundred taken in is therefore not a fund of a hundred to be distributed at will. And the value of the goods on which the service is provided is something else again.

Buying an asset that lasts for years is not like buying a service that is used up at once. Investment has a separate place in the accounts: the machine loses value over time, must be maintained, must be paid for. The sixty in the example are before this wear. To arrive at a net profit you need further information.

Quantity handled, revenue and value added describe three different things: the first a physical activity, the second a transaction, the third the contribution to output. Volume can grow while the unit price falls, and then quantity and value in euros diverge. The model's growth paths are real, that is, adjusted for price changes.

Value added pays labor and property, and contributes to public resources. Labor income becomes consumption, saving and taxes. Capital income does the same, but goes to different people, who spend differently. To follow it, the model builds three groups in each region: the half with the lowest incomes, the middle forty percent, the richest ten percent. They are groups of people ranked by income, not three trades, three generations or three neighborhoods.

Who Owns What

The next table describes the model's starting point. The wealth shares are decisive, because they determine who receives part of capital income, and therefore how the same innovation reaches the different groups. The percentages are shares of income or wealth.

Table 20 – Initial distribution. In each cell: bottom, middle, top group. Anchors: World Inequality Report 2022; some items are model assumptions.

Region	Income shares	Wealth shares
ADV	17%; 43%; 40%	4%; 34%; 62%
EME	12%; 40%; 48%	3%; 27%; 70%
LIC	9%; 36%; 55%	2%; 23%; 75%

In the advanced economies the poorer half receives 17 percent of income, the middle forty percent 43, the richest ten percent 40. But wealth is far more concentrated: the poorer half has 4 percent, the middle 34, the top 62. Half the population owns four cents of every euro.

The anchors come from the World Inequality Report 2022. Income in the advanced economies combines Europe, where the richest ten percent receive 36 percent, and North America. That of the emerging economies uses as references East Asia, where the richest ten percent receive 43, and Latin America, where they reach 55. For low-income countries, the report indicates less than 10 percent of income for the poorer half in sub-Saharan Africa. For wealth in the advanced economies, the European reference is 4, 38 and 58, and the report notes shares of 60–80 percent for the richest ten percent across different areas. The wealth distributions of the emerging and poorer economies lie within the report's ranges.

The groups' shares of labor income are derived so that wages plus capital income, distributed according to wealth, reproduce the overall income shares. It is the way the initial distribution is tied to the model's accounts.

How much does each group spend? In the advanced economies the propensities to consume are 0.85 for the poorer half, 0.58 for the middle and 0.40 for the top. The anchor is a 2020 study by

Fisher and co-authors, with American quintiles mapped onto the model's three groups. The 0.85 means that group spends 85 percent of its disposable income on consumption.

Putting the pieces together, the mechanism becomes clear. If innovation mainly increases income linked to ownership, and ownership is this concentrated, the benefits are concentrated too.

The Owner Who Lives Elsewhere

The place where a good is produced is not necessarily where the income from ownership is collected. In the model, a share ζ of the additional capital income generated by AI in the emerging and poorer economies, net of the hypothetical levy, is paid to the richest ten percent of the advanced economies. It is a rule of geographical attribution, and its strength ranges between 20 and 60 percent.

Not every payment abroad is "leakage". Buying a service, paying a dividend and paying a rent are different things. The model represents a specific channel: that of AI capital income.

Statisticians have a name for this distinction. Gross national income, according to the OECD definition, starts from GDP and adds the balance of labor and property income exchanged with the rest of the world. Payments to non-resident owners can make domestic product and residents' income diverge, as in Ireland. And the distinction concerns where one is economically resident, not one's passport. For example, a small economy produces sixty, receives no income from abroad and pays twelve to foreign owners: its national income is forty-eight. Those twelve remain domestic output, but they are not income for the households who live there.

On a single street the account is even more complicated. Someone living in a neighborhood may receive a salary from a firm in another city, a pension, rent from far away. Someone working

there may go home elsewhere in the evening. Even the bar's purchases have local and distant suppliers.

The model compresses all this into three large regions. The rule on international rents does, however, have a counterpart: what one region pays, the other receives. It allows a distributional channel to be explored without reconstructing every owner. And widening the boundary from the territory to the world changes the meaning of a payment: in the world total, a transfer between regions creates no income. It changes who receives it, and so it can change consumption, saving and investment, even though the sum stays the same.

Three Economies, Three Starting Points

The advanced, emerging and low-income economies bring together populations with very different growth, exposure to technology and resources. Before comparing them you need to know where they start from; pensions and debt return in the next chapter. And watch the weights: GDP shares use market exchange rates, demographic shares use the adult population, and they are not interchangeable.

The advanced economies produce 57 percent of world GDP, the emerging economies 40, the low-income economies 3. The source is the IMF's World Economic Outlook, at market exchange rates.

But the adult population is divided almost the other way round: 14 percent in the advanced economies, 70 in the emerging ones, 16 in the low-income countries. The source is the UN's World Population Prospects 2024.

The share of income going to labor is 55.4 percent in the advanced economies, 50 in the emerging ones, 38.5 in the low-income countries. The source is the ILO; the value for the emerg-

ing economies is chosen so that the world share comes back to 52.4.

Initial exposure to AI is 60 percent in the advanced economies, 40 in the emerging ones, 26 in the low-income countries, according to Cazzaniga and co-authors (2024). Exposure does not mean job loss.

Baseline real annual growth is 1.6 percent in the advanced economies, 4 in the emerging ones, 4.5 in the low-income countries, according to the IMF's medium-term outlook.

Labor supply, between 2026 and 2046, falls in the advanced economies, from minus 0.2 to minus 0.4 percent a year; grows in the emerging ones, from plus 0.6 to plus 0.2; and grows strongly in the poorer countries, from plus 2.6 to plus 2.3. It is derived from the population aged 20 to 64 in the UN projections.

Public revenue is 36.3 percent of GDP in the advanced economies, 26 in the emerging ones, 15 in the poorer countries: IMF DataMapper data for the first, Fiscal Monitor approximations for the others. Gross public debt in 2025 is 108 percent of GDP in the advanced economies, 75.3 in the emerging ones, 48.2 in the poorer countries, according to the April 2026 Fiscal Monitor. The primary balance is negative everywhere: minus 1.6, minus 3.5 and minus 0.8 percent of GDP, derived from the Fiscal Monitor's overall balance minus assumed interest. Public pensions cost 8.1 percent of GDP in the advanced economies, according to the OECD's Pensions at a Glance 2025, and by assumption 4 in the emerging economies and 1.2 in the poorer ones.

One figure among these deserves a second look: the region that produces 57 percent of the world's output is home to 14 percent of its adults. An average weighted by output and one weighted by people describe different worlds. That is why world GDP growth cannot be turned into an improvement for the average household

by multiplying or dividing a single figure. You need the distribution within each region and the weights that connect the regions.

An example shows how much the weighting matters. Two regions: the first has ten adults and produces eighty, the second has ninety adults and produces twenty. The first grows by 10 percent, the second stands still. The total goes from a hundred to a hundred and eight: plus 8 percent. But if each adult is weighted equally, average growth is 1 percent. The first calculation measures how much output has grown; the second gives every person the same weight. Both are correct, and they answer different questions.

For work, the problem is even closer. A region with many adults may have few participants, many young people or many retirees. To obtain the unemployment rate of an aggregate, each region would have to be weighted by its labor force. Every aggregate requires three pieces of information: what is counted, who is included, how much each one weighs.

What remains is what the model does not see. Well-being also depends on the time needed to care for a family member, reach a doctor, follow a course. Having an income does not mean finding a service within reach. The model measures output, income, consumption, employment and public finances; how many doctors there are, how much each home costs and what time the buses run are taken up in chapters 28 and 30.

That is why the productivity of an activity on its own says little. If machines make a service cheaper, a household can benefit without its income rising by a single euro. If an essential service becomes harder to reach, a higher income may not be enough to compensate. The parameter ω counts the benefits of AI services that consumption does not capture, between zero and 5 percent of consumption. Chapter 14 shows its effect, and chapter 28 returns to the concrete case of health care.

And automating one phase is not the whole economy. Manual and relational services can continue to require people: maintenance, care, assistance. But their existence does not guarantee that the new jobs are enough, or within reach of those who lost theirs. Compensation requires quantities, timing, skills, wages, access. It is not enough to list the trades that hold out: you have to see how many people can get into them, at what pay, after how long a wait. A necessary activity may go unfunded, or be done unpaid within families: a social need becomes a paid job only if someone, a firm or the state, pays its cost.

And that is where the next chapter leads. If the division of income between labor and ownership changes, we need to see how the model's tax rules turn that income into public revenue, and how that revenue meets spending needs.

Chapter 11

Public Revenue and Debt

In February 2017, in an interview, Bill Gates proposed something that sounded like a provocation: if a robot does a person's work, the robot should pay taxes too. Just in those days the European Parliament, voting on a report on rules for robotics, had rejected the idea of taxing robot owners to fund training for people who lost their jobs. Since then the question has never gone away: if machines do a growing share of the work, who pays the taxes that workers used to pay?

Gates suggested taxing a robot's work at roughly the level of the person it replaced, and using the money to fund jobs in care and assistance. The idea had a merit: it pointed to the problem. And a flaw: it took for granted that a robot could be counted the way an employee is counted.

This chapter takes the question seriously, but from another angle. It does not talk about robots to be taxed: it talks about income that changes route. A technology can increase output while it becomes harder to fund public services. The result depends on how income is composed and on the rules for revenue and spending. To see the mechanism the chapter uses the rules of an imaginary country.

A Euro That Changes Route

Suppose that in an unidentified country the effective levy is 28 percent on labor income and 19 percent on capital income. They are average values on aggregate bases, contributions included. Under these rules the same total income produces different revenue depending on how it is composed.

Take a total of a hundred: sixty from labor, forty from capital. Revenue is 16.8 from labor and 7.6 from capital: 24.4 in all. Now ten units move from labor to capital. The total stays at a hundred, but the two bases become fifty and fifty. Revenue falls to fourteen

and 9.5: 23.5 in all. The state has lost 0.9, that is, the nine points of difference between the two rates applied only to the part that changed route. Nine cents for every euro that moves from wages to profits. Nobody evaded anything: the income has simply changed form.

But output can grow. If output rises from a hundred to a hundred and ten and remains split half and half between labor and capital, each base reaches fifty-five, and with the same rates revenue becomes 25.85: more than the initial 24.4. Here growth offsets the change in composition. The same technological change can therefore bring lower or higher revenue, depending on how much output it adds and how it distributes it. That is why the full calculation is needed.

And in the full calculation, for the simulated advanced economies, the gain in capital taxes does not offset the loss in labor taxes. Meanwhile, during the transition, some spending on unemployment rises. Taxable income does not disappear: its composition changes, and the need for protection grows.

The Account in Words

Before the formulas, the account in words. The primary balance is what is left of public revenue after spending, before paying interest on the debt. A surplus helps reduce debt, a deficit increases it, and interest on existing debt is added to the bill. In the ratio of debt to GDP, debt is on top and output underneath: growing GDP can slow the rise in the ratio, weak GDP can accelerate it. Both quantities must be followed.

In the model, the public account adds taxes on labor, capital and consumption and subtracts government purchases, pensions, benefits and additional policies: that is the primary balance. The following year's debt adds interest on the existing debt and sub-

tracts the surplus, or adds the deficit. Section A8 of the appendix lists all the items. And raising GDP does not guarantee a better balance: it depends on the tax bases, the rates and spending. The fiscal result must be read together with the division between labor and capital, not only with growth.

A one-year exercise helps avoid confusing deficit and debt. The deficit is a flow during the year; the debt is a stock that contains everything accumulated before. You cannot derive the debt-to-output ratio from a deficit. In an imaginary economy, initial GDP and debt are both a hundred, interest is three, the primary balance zero. With no other operations, final debt is a hundred and three. If final GDP is a hundred and two, the ratio is about 100.98 percent; if it is a hundred and five, about 98.10. Debt has grown by three in both cases, but its weight rose in the first and fell in the second. It was the denominator that changed the reading. The numbers are all in the same unit, and the changes in GDP are nominal, so they can be compared with interest without mixing prices and quantities.

Add a primary surplus of two. Final debt would be a hundred and one, and with GDP of a hundred and two the ratio would fall to about 99.02 percent. The loan still costs money, but surplus and growth offset the interest. The debt relationship in section A8 formalizes this sequence: interest and growth act on the weight of past debt, the primary balance adds or subtracts. To use it, rates and quantities must be kept consistent: putting inflation only in the denominator would distort the comparison without saying how the cost of financing changes.

The International Monetary Fund lists exactly these ingredients for judging whether debt is sustainable: primary balance, real growth, real interest and the level of debt. The exercise is meant to show the mechanism, not to judge a country: the currency of the debt, maturities, refinancing and operations that change debt

without passing through the deficit add further conditions. And the starting level matters. If initial debt is a hundred and fifty instead of a hundred and the rate stays at 3 percent, interest goes from three to 4.5: the same rate on a larger debt costs more.

Over twenty years, the account restarts each year from updated debt and output, while wages, revenue, benefits and pensions change. The effects accumulate, and a temporary change in the balance can leave a mark on debt even after it has passed. And mind the words: one point of GDP of extra deficit for one year is a flow in that year; one point of GDP of extra debt is a difference in the stock, measured against output at that date. They have the same name, but they are not the same thing.

The Debt That Was Already There

Gross debt in the advanced economies in 2025 is calibrated at 108 percent of GDP. Without AI, in the model, it reaches about 132 percent in 2046. Pensions and interest explain most of this underlying pressure: ageing, not technology. The wave of AI is added to a trajectory that was already rising.

That is why the comparison requires two debt paths that change over time. Using a fixed debt in the world without AI would attribute to the technology changes that would happen anyway. The pairs use the same baseline parameters: the difference measures the effect of AI.

In 2036 the primary balance of the advanced economies worsens by about one point of GDP in the median. In 2046 the debt-to-GDP ratio is 11.2 points higher, with a range of 3.4 to 26 points between the best and the worst future out of twenty. Those 11.2 points are the difference between two ratios, not the level of debt.

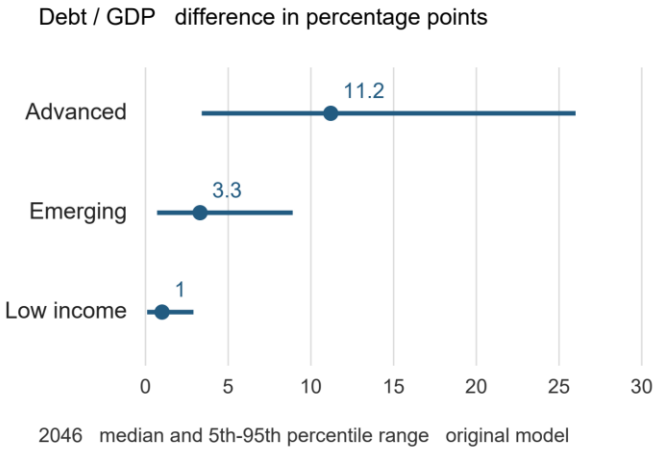


Figure 20 – Additional effect of AI on the gross government debt-to-GDP ratio in 2046. Median and 5th–95th percentile range, in percentage points, for the three regions. Model simulations, 4,000 paired draws.

In the other regions the effect is smaller: 3.3 points in the emerging economies, with a range from 0.7 to 8.9, and 1 point in the low-income countries, from 0.1 to 2.9. But a smaller effect does not mean more capacity to fund services: starting levels, revenue and the model's constraints are different. A poor country with revenue at 15 percent of GDP has less room than a rich one, even if its debt grows less.

An increase in capital revenue, on its own, does not even say whether the public balance improves or worsens. All the items must be read.

Table 21 – Median changes with AI relative to the no-AI path in 2036. Revenue and primary balance are expressed in percentage points of GDP; in brackets, the 5th–95th percentile where available. For LIC only the medians of these three items are reported. Model simulations.

Item	ADV	EME	LIC
Taxes on labor	−1.26 (−2.98; −0.35)	−0.19 (−0.65; −0.02)	−0.01
Taxes on capital	+0.84 (+0.23; +1.98)	+0.19 (+0.02; +0.65)	+0.01
Primary balance	−0.99 (−2.27; −0.28)	−0.22 (−0.76; −0.02)	−0.02

In the advanced economies the median gain in capital taxes is about 0.84 points of GDP; the loss in labor taxes about 1.26. Then the other taxes and other spending come in. And the sum of the medians of the individual items differs from the median of the total: the budget balances in each simulation, before the results are summarized.

Why? An example with three simulations and two items, using numbers without units so they are not confused with the analysis. In the first, the items are zero and a hundred; in the second, one and one; in the third, a hundred and zero. The median of the first item is one, that of the second is one: added together they make two. But the totals of the three simulations are a hundred, two and a hundred, and their median is a hundred. A budget can balance perfectly in every case while a page of medians forms the budget of no case at all. The median is a position in the line, and different cases can occupy that position for different items. The mean of a sum, by contrast, with the same sample and the same weights, equals the sum of the means: the two tools have different properties. That is why the book first calculates each simulated economy and then summarizes.

How Much Would Need to Be Collected

Before changing a tax rate you have to say which base it applies to. A levy on additional income alone and one on all capital income are different experiments.

In 2036 about 4.5 points of GDP move from labor to capital. How much would that portion alone have to be taxed for the primary balance not to worsen? The calculation gives 41 percent, instead of the model's initial 19. It is more than the 28 applied to labor, because the account must also cover the extra benefits.

And a rate that balances the account does not say how much could really be collected: capital moves, income concentrates. It is also different from the five-point increase in the tax on all capital simulated in chapter 13, which redistributes the revenue. And this is where Gates's old proposal shows its limit: the problem is not counting robots, but following income.

To check that the starting path is reasonable, the path without AI was compared with the IMF's projections to 2031.

Table 22 – Gross debt / GDP in 2031. Model without AI at central parameters and IMF Fiscal Monitor projections, April 2026. The revenue rules are assumptions for unidentified countries and do not describe national systems.

Region	Model / IMF	Difference
ADV	113.6% / 114.8%	−1.2 points
EME	85.8% / 88.9%	−3.1 points
LIC	46.9% / 44.0%	+2.9 points

In the advanced economies the model gives debt of 113.6 percent of GDP, the IMF 114.8: a difference of 1.2 points. In the emerging economies the difference is 3.1 points, in the low-income countries 2.9 points in the other direction. The IMF projections were not used to calibrate the model: that is why the comparison is meaningful.

The accounting check asks something else: that every saving, every business result, every public balance and every payment between regions has a counterpart. The model checks household saving, business results, the public balance and payments between regions together, and compares them with output and spending. That way no sum can appear or disappear through an error. The identity and its items are in section A9 of the appendix. Across all simulations the largest discrepancy is 5×10^{-16} of GDP, that is, zero down to the computer's precision.

Chapter 26 feeds the age module's unemployment back into wages and benefits and recalculates debt: the result changes little, from 11.19 to 11.39 points in 2046.

The next chapter opens the box: the assumptions on which all these conclusions depend, and how anyone can change them.

Chapter 12

Calibration, Parameters and Reproducibility

In 2010 two Harvard economists, Carmen Reinhart and Kenneth Rogoff, published a study that seemed to settle a debate: when public debt exceeds 90 percent of GDP, growth collapses. The study was cited by governments and commissions across half the world, in the years of austerity. Three years later a doctoral student at the University of Massachusetts, Thomas Herndon, tried to redo the calculations for a class assignment. They did not add up. With his professors Michael Ash and Robert Pollin he discovered, among other things, that a spreadsheet formula left out some countries. Once the errors were corrected, average growth for countries with debt above 90 percent was no longer negative, minus 0.1 percent, but positive, about plus 2.2: the collapse disappeared.

The lesson is not that economists make mistakes. It is that a result is worth as much as the possibility of checking it. This chapter opens up the book's model: it shows all the assumptions, where they come from, what happens when they are pushed in the direction most favorable to AI, and how anyone can redo the calculations.

When a parameter is uncertain, choosing a single value can hide a decisive part of the problem. That is why the model does not choose a number: it chooses a range, and draws at random within it.

Each uncertain parameter is drawn from a uniform distribution within stated limits. Uniform means the draw does not favor any part of the range.

And the draws are independent: the model imposes no links between one parameter and another. Rapid adoption might come with investment in training, or precede it: later on, an experiment tests exactly these links.

The Twenty-One Assumptions

Here are the model's twenty-one uncertain assumptions, with their range and the anchor used to choose it. The anchors are the reference points in the research. There is no need to memorize them: consult them following the question you are interested in.

The first six concern how much is automated, and when.

1. θ, the automated share of exposed tasks: from 0.25 to 0.70. Anchor: Eloundou and co-authors (2024) find exposure ranging from 1.8 to 46 percent of jobs, depending on whether only the language model or also the software accompanying it is considered; McKinsey (2023) indicates half of all activities by 2045 as an upper bound.

2. φ, the part of exposure that can actually be substituted: from 0.40 to 0.70. Anchor: the IMF estimates 33 points out of 60 in the advanced economies, that is, about 0.55.

3. The central year of adoption in the advanced economies: between 2030 and 2040. An assumption informed by Comin and Hobijn's research on technology diffusion (2010).

4. k, the slope of the adoption curve: from 0.25 to 0.60. An assumption: it governs how fast adoption grows around the central year.

5. The lag for emerging economies: from 1 to 5 years. An assumption, with reference to Comin and Hobijn.

6. The lag for low-income countries: from 3 to 10 years. An assumption, with the same reference.

The regional lags distinguish when a technology exists from when it actually spreads. But for those who lose a task, something else matters too: when the next task arrives. That is why the parameters on new tasks come after adoption.

7. ψ, new tasks for each task displaced: from 0.25 to 0.85. This is the model's most delicate assumption. Acemoglu and Restrepo

(2019) measure a ratio of about 1 between 1947 and 1987 (0.47 of new tasks against 0.48 displaced) and about 0.5 between 1987 and 2017 (0.35 against 0.70). The range follows recent experience and excludes full compensation. Later you will see what happens if it is allowed.

8. The delay of new tasks: from 4 to 10 years. An assumption based on the productivity "J-curve" (Brynjolfsson, Rock and Syverson, 2021) and on the slow appearance of new occupations (Autor and co-authors, 2024).

9. d_0, the annual erosion of the labor share not due to AI: from 0.001 to 0.0035, that is, from 0.10 to 0.35 percent a year. Anchor: ILO data 2004–2024 and Acemoglu and Restrepo 1987–2017. The parameter is compared with these same references in the check further on.

Then come productivity and wages.

10. π, the saving on automated tasks: from 15 to 45 percent. Anchor: Acemoglu (2025) estimates an average saving of 27.

Saving on costs and raising wages are two different steps. How much of the gain reaches workers depends on distribution and bargaining.

11. μ, the gain in tasks where AI helps without replacing: from 0 to 0.30. Anchor: the call center study by Brynjolfsson, Li and Raymond, plus 14 percent on average and plus 34 for the least experienced.

12. β, how much productivity passes into wages: from 0.30 to 1. An assumption.

13. $\varkappa$, how much wages respond to unemployment (the slope of the wage Phillips curve): from 0.05 to 0.50. Anchor: a San Francisco Fed Economic Letter (2017–30) estimates 0.45 between 1991 and 2008 and 0.14 between 2009 and 2015; a Bank of Canada study (2019–8) finds the same flattening after 2008 in ten advanced economies.

14. s, the share of AI rents going to workers: from 0 to 0.30. An assumption of the unchanged-policy scenario: bargaining and profit sharing.

Part of income becomes consumption, another part saving and investment, and some rents end up outside the region. To know whether output finds buyers, the model must also represent these responses.

15. η, how much investment responds to the profit share: from 0 to 1. An assumption, linked to the debate between those who think demand is wage-led and those who think it is profit-led.

16. ν, how much a fall in demand passes into output, after the offsetting effects of monetary and fiscal policy: from 0.30 to 1. An assumption.

17. ζ, the share of AI rents in emerging and poorer economies booked abroad: from 0.20 to 0.60. An assumption, linked to the concentration of AI firms.

18. x, how much exposure grows by 2046: from 1 to 1.60, that is, up to 60 percent more. An assumption based on expanding capabilities and on growth in the robot stock of 9 percent a year (IFR, 2025). And it is this 60 percent limit that excludes the extreme scenario, that of an AI able to do almost everything.

The last three choices answer three different questions: benefits that money income does not capture, the cost of interest, and the response of labor demand to wages.

19. ω, the well-being AI services provide without passing through consumption: from 0 to 5 percent of consumption. It is an assumption used only in the welfare-adjusted results, never in the main measures.

20. Δr, the change in the real interest rate: from minus one point to plus one percentage point. Note: one point, not 1 percent of the rate.

21. ε, how much labor demand responds to wages: from 0.15 to 0.75. Anchor: Hamermesh (1993), with a reference value of 0.30 and the same range; the meta-analysis by Lichter, Peichl and Siegloch (2015) finds an average of about 0.25. This parameter was added with version 2 of the model.

In the code the names are, in the same order: theta, phi_s, mid_adv, k_slope, lag_eme, lag_lic, psi, lag_new, d0, pi_save, mu, beta_w, kappa, rent_share, eta, nu, zeta, exp_mult, omega, dr, eps. The table is generated by the file calibration/data.py, which also contains the regional and distributional constants of chapter 10.

The Bag of Possibilities

The method the book uses to explore four thousand futures has a curious history. In 1946 the mathematician Stanisław Ulam, convalescing after an illness, was playing solitaire and wondered what the probability of winning a game was. Calculating it with formulas was impossible. He had a simpler idea: play a great many games and count how many are won. With John von Neumann he turned the idea into a method of calculation for problems in nuclear physics at Los Alamos. The name, inspired by the Monte Carlo casino that an uncle of Ulam's frequented, is usually attributed to his colleague Nicholas Metropolis.

The Monte Carlo method draws at random, but it draws from possibilities chosen beforehand. Before the draw someone decides which values are allowed and how much each weighs. A computer can make millions of draws, but it always works on the distributions it has been given.

Imagine a bag with ten tickets: six say "loss", four "gain". If you draw many times, putting the ticket back each time, losses will tend towards 60 percent. With nine "loss" tickets and one "gain", they will tend towards 90. No speed of calculation tells you which

bag describes the world. What goes into the bag needs to be justified.

In the model the tickets do not show an outcome: they contain parameter values. The equations turn each combination into a trajectory. But the point remains: the frequencies of the results depend on the structure of the model and on the weight given to the initial combinations. When the book says that the poorer half loses in 92 percent of the simulations, it means it almost always loses with the tickets the model put in the bag. The bag is here, open, and anyone can change it.

The same method is used in many fields to carry uncertainty from inputs to results. The NIST Uncertainty Machine, the American standards institute's tool for measurement uncertainty, asks for three things: the distributions of the inputs, a function linking them to the result and, when needed, the correlations between inputs. It is the same method.

Now take two questions with two answers each: is adoption fast or slow? Is training accessible or difficult? There are four combinations. If each answer weighs a half and the questions are independent, each combination weighs a quarter.

Now imagine a different world, in which half the cases combine fast adoption and accessible training, and the other half slow adoption and difficult training. The individual answers still have the same weight: half fast and half slow, half accessible and half difficult. But the combinations have changed. And a result that depends on the meeting of speed and access can change even if every single range stays identical.

Ranges, then, are not enough to describe a model: how the parameters are combined also matters. The model uses independent draws, and the second experiment in this chapter tests what changes with a link.

There is also a question of time, which the bag does not show. The delay of new tasks represents a period in which income and experience may be missing. What matters is how long it lasts, and how long people can hold out while waiting.

But be careful not to confuse scales. A delay of four to ten years does not say how long it will take a person to find work. It says when the model makes new tasks available relative to displacement. The age module then adds other obstacles to matching people and jobs.

And ψ, between 0.25 and 0.85, deserves attention: by construction it excludes every lost task being replaced by a new one. The choice follows recent data. The next experiment tests what happens if it is removed.

Three Objections, Three Tests

It is the optimists' first objection, and it is legitimate: the model is pessimistic because it excludes new tasks fully offsetting those lost. So let us remove that limit.

The experiment holds all other draws fixed and widens ψ from 0.25–0.85 to 0.25–1.20: up to more than one new task for each task displaced. It uses four thousand pairs of simulations, with and without AI. ψ says how many new tasks appear; the delay still decides when. Allowing ψ equal to or above one opens the door to full compensation. It does not make it arrive in the same year as displacement.

Table 23 – Experiment on new tasks. Medians, AI minus no AI. World bottom-half consumption: percentages; ADV unemployment: percentage points. First row: widened ψ range; second: subset with ψ of at least one. The other parameters keep their own design; the subset is not a new estimated population.

Experiment	Consumption 2036 %	Consumption 2046 %	Unempl. 2046 pp
ψ between 0.25 and 1.20	−0.97	−2.15	+0.15

Experiment	Consumption 2036 %	Consumption 2046 %	Unempl. 2046 pp
Only ψ of at least 1	−0.80	+1.67	+0.00

With the wider range, the median loss of the poorer half in 2046 shrinks compared with chapter 14, but remains negative. Additional unemployment in the advanced economies approaches zero. A stronger recovery changes the picture, with timing and distributional effects still to be followed.

Then let us look only at the most optimistic futures, those in which new tasks are at least as many as those lost. There are 857 combinations, about 21.42 percent. Here, at last, median consumption of the poorer half in 2046 is positive: plus 1.67 percent. Additional unemployment, in the median, is zero.

But look at the rest. Between the worst future out of twenty and the best out of twenty, consumption ranges from minus 5.82 to plus 9.45 percent. About 35.59 percent of these optimistic futures still show a loss. For unemployment, the worst future out of twenty reaches plus 3.32 points. And the zero of the median is rounded: it does not mean zero in every simulation.

And there is time. In 2036, even in these optimistic futures, median consumption of the poorer half is still negative. New tasks can be enough and arrive too late. Meanwhile lower wages, the distribution of rents and the underlying erosion of the labor share remain.

This test says something precise: if new human tasks were numerous and accessible enough, the loss of the poorest would shrink a great deal. That is therefore the figure to watch. And even in the most favorable futures, a third still have a loss.

Second objection: what if those who adopt AI first are also those who prepare first? The second experiment answers the question about links between parameters.

It keeps the same ranges for adoption, delay and ψ as the original model, and changes the way they are combined. In "coordinated preparation" early adoption comes with short delays and more new tasks. In "late preparation" it comes with long delays and few new tasks. A control case combines the three independently.

For those who want the technical detail: the link is built with a Gaussian copula. Normal variables are correlated and transformed into percentiles of the ranges. The correlations use an absolute value of 0.5. The other distributions remain independent, and the same draws, different from the initial sample, are used to compare the three cases.

Table 24 – Adoption and preparation paired: medians of the differences AI minus no AI, 4,000 pairs per configuration. Consumption: world bottom half, percentages. Unemployment: ADV, percentage points. The independent check uses the same normal draws as the other two cases.

Combinations	**Consumption 2046 %**	**Unempl. 2036 pp**	**Unempl. 2046 pp**
Coordinated preparation	−3.14	+4.02	+1.28
Independence	−3.34	+4.30	+1.37
Late preparation	−3.50	+4.51	+1.26

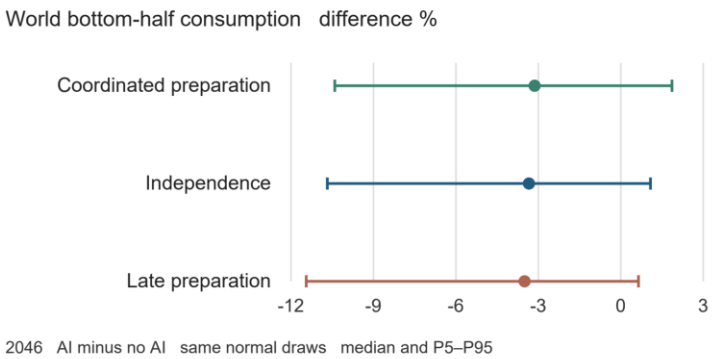


Figure 21 – World bottom-half consumption in 2046: AI minus no AI. Median and P5–P95 in the three pairings of adoption, lag and reinstatement. The marginal ranges remain the original ones; the assumed dependence changes. Same normal draws in the three comparisons.

Preparing in time helps: it reduces the loss in consumption in 2046 and the jump in unemployment in 2036. Preparing late worsens both. But in all three cases median consumption of the poorer half remains negative.

And in 2046 unemployment does not follow a simple ranking: plus 1.28 points in the coordinated case, plus 1.26 in the late one, plus 1.37 in the independent control. Changing the timing changes the path, but not the conclusion.

To know which link holds in reality you need data linking adoption, training, new tasks and youth entry: that is the work proposed by the protocol in chapter 18.

Third objection: perhaps the model underestimates productivity gains. More productivity increases output and softens some losses; to know who benefits, the distributional account must be redone.

In 2036 world GDP with AI is 1.3 percent higher in the median, with a range between the worst and the best future out of twenty of 0.3 to 3.5 percent. The limit estimated by Acemoglu (2025), 0.66 percent of total factor productivity over ten years,

corresponds to about 0.9 percent of GDP. The model's median therefore lies in the cautious part of the literature; higher estimates require faster and broader adoption. So let us look at the most favorable futures.

Table 26 – Checks in favorable subsets. Changes relative to the paired no-AI simulation; consumption as a measure of purchasing power.

Selection of simulations	**Bottom group, 2036**	**Other results**
Fifth with the largest cost savings	Median −0.6%; worse in 82% of draws	In 2046 median −2.1% for the high-savings simulations
Savings and gains in assisted tasks both in the top fifth	Median −0.3%; worse in 64% of draws	The top 10% gains 3.8%

In the fifth of futures with the highest cost savings, the poorer half still loses in 82 percent of cases, with a median of minus 0.6 percent in 2036 and minus 2.1 in 2046. When savings and gains from assistance are both in the top fifth, the median loss falls to minus 0.3 percent, but the poorest still lose in 64 percent of cases. And the richest ten percent gain 3.8.

More productivity softens the losses of the poorest. It does not change who takes the gain. The problem is not that AI produces too little: it is who gets what it produces.

And ω, the well-being that does not pass through consumption, should be looked at next to the purchasing power tables, not inside them. They are two different questions: how much money a person receives, and what additional services they can use.

Three tests, and one warning that applies to all of them: more simulations reduce sampling noise, but they do not correct a forgotten parameter or a wrong relationship.

The bands in the charts, between the fifth and the ninety-fifth percentile, describe the central 90 percent of the futures simulated.

And there is a subtler trap. A curve that takes the median of each year can join points that come from different simulations. Take three cases and three years. The first goes minus one, plus one, plus one. The second plus one, minus one, plus one. The third plus one, plus one, minus one. Every year the median is positive. Yet each of the three had a bad year. The median chart draws a road that always climbs, which nobody actually travelled.

Following a whole path is different from looking at one year at a time. For people, you need data that follow them over time; for the model, the identity of each simulation from start to finish. Duration, the moment of entry, cumulative recovery: they all require the whole path. That is why chapter 14 also looks at the unemployment peaks of each individual future.

The Past as a Test

A model about the future must at least not contradict the past. But a figure used to choose a parameter must be distinguished from a figure kept aside to check it. The first comparison concerns the fall in the labor share that took place before generative AI.

Table 25 – Historical comparisons. The third indicator is a percentage change in labor content; the first two are percentage points of a share.

Reference	Observed data	Model range
ILO, global labor income share, 2004–2024	−1.6 percentage points	−1.1 to −3.7 points
Karabarbounis and Neiman (2014), labor share in global firms, 1975–2012	−5.0 percentage points	−2.3 to −7.8 points
Acemoglu and Restrepo (2019), labor content of tasks in the US, 1987–2017	About −10%	Not available

The annual erosion implied by these declines is 0.15, 0.22 and 0.35 percent. All three lie within the range chosen for d_0; the last

coincides with the upper limit. The global declines lie in the lower half of the model's range: using only American data would overstate the global decline.

These same references were used to choose the range for d_0: the comparison is a consistency check, not independent proof.

The second check is a piece of history. Friedrich Engels, in 1845, described the condition of the workers of Manchester: the wealth of the factories grew before their eyes, and their lives did not. The economic historian Robert Allen (2009) measured that period: in Britain between 1800 and 1830 output per worker grew by 0.63 percent a year, while real wages stood still.

In the model's task structure, something like this requires net displacement of about 0.6 percent a year, with very few new tasks. The labor share would fall by about 17 percent in thirty years. Low values of ψ with rapid adoption can produce displacement of this size in the advanced economies. And it is no museum piece: in American manufacturing between 1987 and 2017 displacement was about 1.1 percent a year.

The model can reproduce Engels' pause. It is proof that the mechanism is real. Output growing while wages stay flat for a whole generation has happened before.

The simulated wave of AI has not happened yet: no check on the past can verify it.

What the checks do say is something else, and it is important: the book's central result, growth with worse distribution, survives even when the assumptions are pushed in the direction most favorable to AI.

Redoing the Sums Yourself

The code is public. This is the map for anyone who wants to open it.

The file aimodel/calibration/data.py contains the calibration values and the ranges of the twenty-one parameters. The file aimodel/core/economy.py contains the equations: tasks, labor market, households, firms, the state, regions and the accounting check, that is, the relationships in appendix A. The file aimodel/core/run.py manages the paired simulations: the same draw of parameters, with and without the wave of AI. The file aimodel/analysis/metrics.py calculates the result measures: GDP, labor share, consumption by group and region, unemployment, debt. The file aimodel/analysis/montecarlo.py runs the sampling and the Sobol sensitivity analysis. The file aimodel/analysis/backtest.py runs the historical comparisons. The file aimodel/analysis/policy.py contains the nine interventions. The script run_all.py runs everything and saves the results in the outputs folder. Later extensions have their own configurations: chapter 26 distinguishes the experiment on the loop between work and income from the earlier calculations.

You need Python 3.10 or later, with NumPy and SALib. There are three commands: pip install numpy SALib to install, python run_all.py for the full analysis, python run_all.py --fast for the quick version, which skips the Sobol analysis and uses 500 draws for the policies. On the development computer the run takes about six seconds; the time varies from machine to machine.

Table 27 – Files produced by the software. The script also prints the main results, for comparison with the figures in the book.

File in outputs/	Content in the model
headline.json	Aggregate, income-group and fiscal results
backtest.json	Historical comparisons
policy.json	Nine interventions
sobol.json	Sobol sensitivity

The settings are fixed. Monte Carlo: 4,000 paired draws, seed 20261010, independent uniform distributions. Policies: 2,000 paired draws per intervention, with the same seed and the same draws as the unchanged-policy case. Sobol: Saltelli sampling with N equal to 1,024, that is, 23,552 paired simulations, seed 7, first-order and total indices. Horizon: from 2026 to 2046, year by year. Accounting check: sectoral balances must sum to zero in every year of every simulation, with a stated maximum violation of 5×10^{-16} of GDP, a computer rounding error.

The results were regenerated with the software before the text was finalized and match the figures in the book.

The seed is the starting point of the random number generator. With the same seed, the same algorithm and the same sequence of calls, the computer regenerates exactly the same values. The book's project keeps the results and copies of the original data: whoever checks can trace the numbers in a figure back to the files actually used.

Are four thousand draws enough? It depends on what you look at. A median stabilizes quickly; a rare frequency needs more draws; a result close to a threshold requires more care.

And there is a rule that holds the whole book together: the comparison is made within each pair. Take any two pairs. In the first, consumption without AI is 100 and with AI 98. In the second, 150 and 151. The effects are minus two and plus one. Comparing the 98 of the first with the 150 of the second would give minus 52: a number that mixes different starting conditions with the effect of the technology. The same rule applies to policies.

Narrowing or widening a range in data.py shows how much a result depends on an assumption. Adding an entry to policy.py makes it possible to test a new intervention on the same draws. Setting the AI parameter to zero reproduces the path without AI.

Anyone who disagrees with the calibration is invited to change it and publish the result.

It pays to change one thing at a time: changing seed, definitions and policies together makes it impossible to see which change produced the difference.

The age module has its own configuration, in the file config/employment.json, and chapter 26 links it to public finances.

A serious experiment must be able to prove its author wrong. For each test you have to say in advance which result would count against the hypothesis. If every outcome is reinterpreted as confirmation, the experiment tests nothing. The tests in this chapter, widening ψ and linking adoption and preparation, were built this way: they could have refuted the book's thesis. They softened it in some futures. They did not overturn it.

Speeding up retraining and providing income during the wait address different problems. One measure can improve consumption without increasing employment; another can reduce unemployment while increasing public spending.

And before reading the policy results, the comparison must be fixed. You can compare the world with AI and the world without AI, with the same policy. Or a policy with the case in which nothing is done. They are not the same thing. A transfer that is useful even without AI can bring a large benefit compared with a world without that transfer, and still leave a negative effect of AI.

The next chapter compares the interventions, keeping objective, cost and conditions explicit. You will not find a guaranteed recipe. You will find the differences between choices, and what each one costs.

Chapter 13

Public Policy Evaluation

Since 1982, every autumn, almost every resident of Alaska receives a check. They do not have to do anything to get it: it is enough to have lived in the state for at least a year. The money comes from a public fund fed by oil revenues, and the amount changes every year with the fund's returns, in recent decades often between one and two thousand dollars per person. Children included. For forty years critics predicted that an unconditional check would take away the will to work. In 2022 Damon Jones and Ioana Marinescu checked: total employment in Alaska did not fall.

It is one of the few real experiments resembling one of the policies discussed in this chapter: giving everyone a share of the income from common wealth. But supporting income, distributing ownership and creating job opportunities are different things, working through different mechanisms. To compare them you have to say which problem you want to reduce, and against which alternative. A benefit for consumption is not automatically an extra job.

Two Comparisons That Change the Answer

In this chapter's second table, a hypothetical tax on capital raises consumption of the poorer half by 8.2 percent, while the AI dividend stays close to 1. The first seems to win by a mile. But the difference does not prove that the first measure better corrects the damage done by AI: the two numbers contain different amounts of redistribution.

The book uses two comparisons. In the first, the world with AI and a given policy is compared with the world without AI applying the same policy. The question is: given those rules, does introducing AI make people's lives better or worse? In the second, the world with AI and the policy is compared with the world without

AI and without that policy. Here the result also contains the redistribution the intervention would produce anyway, AI or no AI.

Both comparisons make sense, but they answer different questions. A general tax on capital raises the consumption of the poorest because it shifts resources from one group to another, and part of that shift would happen even without automation. The first comparison is written like this, with P the policy, the same in the two simulations, and any outcome, as long as it is defined the same way:

$$\Delta AI(P) = Outcome(AI, P) - Outcome(no\ AI, P)$$

To know instead how much a policy helps in the world with AI, you compare the world with AI and the policy with the world with AI and no measures. It is the useful comparison for choosing between interventions.

The model's interventions begin in 2027. Each is tested on the same 2,000 draws used for the no-measures case, and that is why the medians in this chapter differ slightly from those of the 4,000 draws in the general results.

The Work Still to Be Found

The first table concerns work and output. For 2036, unemployment in the advanced economies is a difference in points and world GDP a percentage change, both relative to the world without AI with the same policy.

Table 28 – Labor and output in the nine interventions and in the unchanged-policy case. Medians of the model's 2,000 paired simulations for the policy comparison. The funded package adds the capital tax and the taxation of rents where AI is used. Model simulations.

Intervention	ADV unempl. pp	World GDP %
Unchanged policies	4.44	1.31
Capital tax	4.42	1.34
AI dividend	4.43	1.33
Income support	4.42	1.34
Training and new tasks	3.30	1.38
Bargaining and profit sharing	4.43	1.32
Delayed adoption	3.49	0.88
Rents taxed where used	4.44	1.32
Package	3.30	1.39
Funded package	3.30	1.39

The comparison of policies uses its own sample, of two thousand pairs of worlds. There, with no measures at all, AI adds 4.44 points of unemployment in the advanced economies in 2036, just above the 4.3 of the four thousand pairs in chapter 6: the difference comes only from the sample. And what is striking is how many rows stay close to that number. Dividend, income support, bargaining: they change resources and their distribution, but they barely touch unemployment. A transfer supports an income without creating a hire.

The chart makes this distinction visible for some interventions. The bars start from zero and show how far unemployment ex-

ceeds that of the respective world without AI. Even the shortest bar stays above zero.

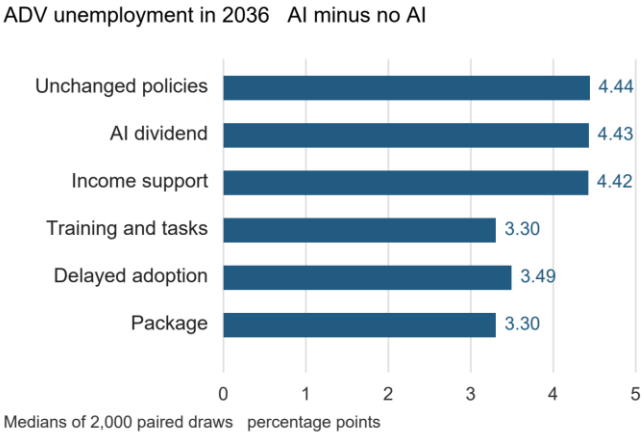


Figure 22 – Additional effect of AI on unemployment in the advanced economies in 2036. Medians from Table 28. The chart shows a selection of interventions; Table 28 keeps them all. Model simulations.

Training and new tasks reduce the increase from 4.44 to 3.30 points: about a quarter of the increase, not a quarter of the unemployed. World GDP grows a little more. Both how many new activities arise and when they arise matter. The empirical reference for this policy is training programs tied to specific sectors, which according to Katz and co-authors (2022) raised participants' earnings by between 11 and 40 percent. The simulated policy links training to faster creation of new tasks across the whole economy.

And even with training, the poorer half still loses compared with the world without AI in 81 percent of the 2036 simulations, and in 2046 it remains 1 percent below. Recovering work does not automatically give back income already lost, nor change how gains are shared. Supporting spending during the wait remains a separate goal.

Distributing the Income from Ownership

The AI dividend works like the Alaska fund, but with a different source: half of the increase in capital income produced by AI goes into a fund, which distributes it in equal shares to every person. The base is the additional increase, not half of all existing profits.

The fund changes who receives the returns. In 2036 the effect of AI on the poorer half becomes about plus 1 percent, positive in 90 percent of the simulations. In 2046 the benefit falls to plus 0.6, while pressure on wages continues.

The direct cost to the state, in the model, is small, because the mechanism moves capital income from one pocket to another. In the real world the fund has to be built: who contributes the shares, how gains due to AI are measured, who governs the rights. Alaska has a fund fed by a publicly owned natural resource; a global fund fed by AI rents is something else. And the fund answers the problem of income: in the chart the increase in unemployment stays almost the same. A dividend does not open the door to those looking for a first job.

Income support unconnected to work looks more directly at those who wait. With the first comparison, for the same measure, the effect of AI on the poorer half in 2036 becomes slightly positive, about plus 0.3 percent; but in about a third of the simulations a loss remains, and in 2046, with unemployment still high, the loss returns to 2 percent.

In the next table the same support shows plus 3.3 percent in 2036 instead, because there the reference is the world without AI and without the policy. It is exactly the earlier distinction: support also shifts resources that have nothing to do with AI.

And a useful transfer has a cost. Support costs about 0.7 percent of GDP a year and, if it is not financed, adds about 11 points

of world debt in 2046 compared with the no-measures case. You have to say who pays.

Table 29 – Public finances and consumption in the same experiments as Table 28. Balance: difference in the average primary balance for 2030–2040 in points of GDP relative to unchanged policies. Debt: difference in the world debt-to-GDP ratio in 2046, in the same comparison. Consumption: percentage change per capita of the bottom half in 2036 relative to the world without AI and without the policy. These are different comparisons, indicated deliberately. Model simulations.

Intervention	Balance pp	Debt pp	Consumption %
Unchanged policies	0	0	−1.1%
Capital tax	+0.35	−6.0	+8.2%
AI dividend	+0.05	−1.0	+1.0%
Income support	−0.71	+11.1	+3.3%
Training and new tasks	−0.01	−1.4	−0.7%
Bargaining and profit sharing	+0.07	−1.4	−0.4%
Delayed adoption	+0.14	−1.2	−0.6%
Rents taxed where used	0.00	−0.1	−1.1%
Package	−0.58	+7.0	+6.1%
Funded package	−0.24	+1.3	+15.1%

On the fear that money takes away the will to work, real evidence is more reassuring than is often said. In Finland, in 2017 and 2018, two thousand unemployed people drawn by lot received 560 euros a month with no conditions: the effect on employment was small, while well-being and trust improved. Alaska did not see work decline. In Kenya cash transfers led to new businesses without reducing hours worked. The prejudice is refuted.

Bargaining, Delay, Revenue

Bargaining and worker profit sharing raise the pay of those who work and slightly improve public finances. But they distribute little to the poorer half, because in the calibration that half receives a small share of labor income. And they do not create jobs. A rule that improves the wages of the employed can do little for those who have yet to get in: they need tools aimed at a different situation.

If instead, in an imaginary country, the levy on capital is raised by five points and the revenue returned to every person, in the model the balance and debt improve, with debt six points lower than in the no-measures case. The plus 8.2 percent in consumption, however, contains redistribution that would happen even without AI. It is a broader lever than the dividend, which only affects additional income.

Delaying adoption by two years reduces the rise in unemployment to 3.49 points in 2036. But the GDP gain falls from 1.31 to 0.88 percent: about a third of that year's benefit is given up. And the poorer half is not protected. The time gained can be very valuable, but delay alone brings no shared ownership, no support and no new opportunities to get in. It creates a window; filling it requires services and interventions people can actually use.

Taxing AI rents in the country where AI is used addresses instead the boundary between economies. It changes the overall result for the poorer half little, but it shifts revenue from the advanced economies towards the emerging and poorer ones. It is the question of chapter 10: where does the value produced or paid stay? An almost unchanged world total hides a very different distribution between territories.

The Measures Together, and the Door Still Closed

The package combines the dividend, income support, training with new tasks, bargaining and profit sharing. Compared with the world without AI with the same policies, the poorer half moves to plus 3 percent in 2036 and plus 3.3 in 2046. Only 0.2 percent of the simulations keep a loss: two in a thousand.

The plus 6.1 percent in table 29 still uses the reference without AI and without the policy, and in that comparison the financed package reaches 15.1. Part of those numbers comes from general redistribution.

Without financing, the package increases world debt by seven points in 2046 compared with the no-measures case. Adding the capital tax and the taxation of rents where AI is used, the increase falls to 1.3 points. It is close to fiscal neutrality; the pressures of ageing remain.

And here is the limit that matters most. Additional unemployment remains at 3.30 points in 2036. The package protects consumption much better, but it does not close the labor gap. None of these tools acts directly on hiring. That is why chapter 24 tests three other tools, which act precisely on hiring: incentives, public employment and reduced working hours.

The limit weighs above all on first entry: none of the nine measures opens the door to those looking for a first job.

There remains the question of time. In half the simulations the peak of unemployment falls between 2036 and 2041: it is a strong reason to build the tools before the phase of greatest pressure, not during it. The conclusion favors tools that change ownership and job creation over simply slowing the technology. Consumption, employment and debt do not always give the same ranking: you have to choose what you want to achieve.

The measures can protect income without closing the gap between labor demand and supply. The needs for care, maintenance and teaching do not become paid jobs by themselves: resources, demand and organization are needed. The next chapter looks at how to read all these results without hiding them inside an average.

Chapter 14

Aggregate Results and Heterogeneity

There is an old joke about statisticians: one of them drowned crossing a river that was, on average, one meter deep. It is funny because it is true. A small average can contain points where you drown. A modest average loss can come with very great difficulties for some people. To read this book's results you need to know who is inside each group, and which internal differences the model is able to see. This chapter puts output, distribution, employment and the benefits of services side by side.

How Much a Small Number Contains

In the model, real consumption per person measures the resources available to buy goods and services. "Real" means that purchasing power is compared, not just a sum in euros. The poorer half is a group defined by income: it is not an age group, a trade or the set of the unemployed. Inside it there are people who keep their jobs and people who lose them.

In 2036 the median difference for this group is about minus 1 percent compared with the world without AI. That does not mean everyone loses exactly 1 percent. The number summarizes the group in each simulation, and the median then summarizes the simulations. Those who lose their jobs and those looking for their first suffer far greater harm than the average.

The next table puts the two worlds on the same page. All results refer to 2036, with unchanged rules, and are differences between paired simulations, with and without AI. ADV indicates the advanced economies, EME the emerging ones. GDP and consumption are percentage changes; shares, unemployment and the debt-to-GDP ratio are points. The last column says in what share of simulations the indicator worsens: for unemployment and debt, worsening means rising; for consumption, output and the labor share, falling.

Table 30 – Two worlds in 2036. Medians, 5th and 95th percentiles and frequency of worsening in the model's 4,000 draws. The row with unmeasured benefits is a measure adjusted with an additional assumption; it remains separate from headline consumption. Model simulations.

Indicator	Median	5th–95th	Worsens
World GDP	+1.3%	+0.3; +3.5%	0%
World labor share	−3.0 pp	−7.6; −0.75 pp	100%
Bottom-half consumption	−1.0%	−4.9; +0.1%	92%
Bottom half with unmeasured benefits	−0.1%	−3.5; +1.9%	55%
Middle 40% consumption	−0.3%	−2.4; +0.6%	70%
Top 10% consumption	+3.1%	+0.7; +7.8%	0%
ADV unemployment	+4.3 pp	+1.4; +10.7 pp	100%
EME unemployment	+1.9 pp	+0.3; +6.0 pp	100%
ADV debt / GDP	+1.5 pp	−0.3; +7.2 pp	88%

Worldwide, output grows by 1.3 percent and consumption of the poorer half falls by 1. The richest ten percent gain 3.1 percent in the median; the middle forty percent lose about 0.3. Total growth and the distribution of benefits take different paths.

Two regularities hold in almost all simulations. Gains concentrate at the top in every draw, and regions benefit asymmetrically in almost all of them. What remains uncertain is how large the damage is: for the poorer half it ranges from a loss of 4.9 percent in the worst future out of twenty to a gain of 0.1 in the best. Even in the luckiest case, the poorer half gains a tenth of a point while GDP grows by more than a point.

In 2046 the consumption gap of the poorer half reaches minus 3.3 percent, with world GDP 3.2 percent higher and the labor share 5.9 points lower. Waiting is not enough: the gap keeps wid-

ening and flattens only after 2044. In 2046 it remains negative in 89 percent of the simulations. Higher output, a lower labor share and consumption divided differently coexist: the first figure does not cancel the others.

The Peak the Average Hides

In the advanced economies, the median trajectory of additional unemployment reaches about 4.9 points in 2039. It stays above four points from 2036 to 2041, and in 2046 it is still at plus 1.4. But each individual simulation has its peak in a different year. If you take the peak of each and calculate their median, you get 6.3 points, with a range from 2.6 to 12.8. It does not coincide with the maximum of the median curve.

These are two different questions. The year-by-year median says what a typical year is like. The median of the peaks says how high a typical path climbs, sooner or later. Since the peaks do not all fall in the same year, the summary curve looks lower than the hardest experience of the paths it summarizes. It is the river that is one meter deep on average.

In six simulations out of ten, unemployment in the advanced economies in 2046 is still more than half a point higher. In the fifth of simulations where wages are most rigid, the final difference is four points. In the emerging economies the increase is 1.9 points in 2036, reaches about 3.3 around 2041 and remains at 2.3 in 2046.

How many people is that? Applying the 2036 differences to labor forces of about 0.6 billion in the advanced economies and 2.5 billion in the emerging ones gives about 26 million additional unemployed in the former and 47 million in the latter. Seventy-three million people, more than the population of the United Kingdom. World population in 2036, according to the UN's 2024 projec-

tions, will be about 8.95 billion. They are orders of magnitude, obtained by applying the model's rates to the labor forces.

Services That Cost Less

Not everything moves in one direction. A translation, help with studying, a first consultation can become much cheaper even for those who have lost income. The book considers advice, translation, tutoring and medical triage among the services AI could make less expensive. The model counts this benefit with a separate parameter, that of the well-being consumption does not measure.

With that adjustment, in 2036 the median loss of the poorer half almost disappears: minus 0.1 percent instead of minus 1. Worsening appears in 55 percent of the simulations instead of 92, with a range from minus 3.5 to plus 1.9.

A useful service can improve life without paying the rent. The benefit must be counted, but only once: if the saving is already included in another measure, it must not be added again. And for the larger gaps of the 2040s the adjustment is not enough to offset the loss. Much of the consequences for job quality, the environment and relationships also remain outside.

The same goes for health. A faster diagnosis does not say who will provide the treatment, how long the wait will be, which care will be within reach. Technical capability spreads slowly through hospitals. Ignoring the benefits because work gets worse would be a mistake; presenting a technical possibility as care that has already arrived would be the opposite mistake.

The next table brings together model results and qualitative assessments, dimension by dimension of daily life. The last column says what each row rests on.

Table 31 – Everyday life in the worlds with and without AI. Comparisons by dimension, with distinct bases and judgments. Year 2036 unless otherwise indicated. Model simulations and evidence cited in the text.

Dimension	Without AI	With AI, and judgment	Basis
Low and middle incomes	Slow real growth; labor share already falling	Output about +1%; labor share −3 pp. Bottom half: consumption −1.0% in 2036, −3.3% in 2046; middle 40% −0.3% in 2036. Mild average deterioration, sharper for displaced workers.	Model
Top 10% and power	Slow growth	Consumption +3.1% world, +5.7% ADV in 2036. Benefit for the group, but greater concentration of income, market power and tax base in a few firms and countries.	Model; distributional judgment
Work	Shortages where the labor force is ageing, including China	World demand −2.1% for the same supply. ADV unemployment +4.3 pp in 2036 (26 million), +1.4 pp in 2046; EME +1.9 pp in 2036 (47 million). Fewer administrative, clerical, junior professional and customer support roles. Persistent and unequal deterioration.	Model and literature
Young people	Entry-level roles present; slow careers	Fewer junior hires in exposed office jobs; first step harder. Judgment: clearly worse.	Evidence; outside the original macro model

Dimension	Without AI	With AI, and judgment	Basis
Ordinary services	Health, legal, education and administrative help expensive and scarce	Routine advice, translation, tutoring, triage and documents cheaper and in more languages. Improvement of unknown size: it could soften the average income loss.	Judgment; effect not measured
Healthcare	Steady progress; shortages of doctors and nurses growing with age	Diagnosis and drug discovery probably faster. Slow diffusion to patients through regulation and health systems. Possible moderate and uncertain improvement.	Judgment
Public finances and welfare	ADV debt from 108% toward 132% of GDP by 2046 because of ageing	About 1% of GDP of ADV primary deficit is added to the same pressure because of the shift from wages to profits. Mild deterioration, unless taxes follow capital income.	Model
Low- and middle-income countries	Growth from a young population and manufacturing	Substitution with local profits partly booked abroad; bottom half EME −1.1%, LIC −0.4% in 2036. Outsourcing at risk; cheap services for farmers, clinics and schools. Mixed balance: better access, more fragile careers.	Model; judgment on services
Energy and environment	Growing electricity demand	Data centers from 415 TWh in 2024 to about 945 in 2030 in the IEA 2025 base case. Higher demand; possible benefits for grids and materials. Mixed balance.	IEA 2025; judgment

Dimension	Without AI	With AI, and judgment	Basis
Older people and care	Shortages growing with the population over 80	Monitoring, assistance and some robotics ease the burden; personal care remains mainly human. Small improvement according to the analysis.	Judgment
Information and trust	Known problems of disinformation	More synthetic content and surveillance capacity, together with better tools against fraud. Judgment: slightly worse, depending on governance.	Judgment

Among these rows, one concerns energy. In the International Energy Agency's base case (2025), data centers go from 415 terawatt-hours in 2024 to about 945 in 2030. A terawatt-hour is a billion kilowatt-hours. It concerns data centers, not all the world's electricity nor AI alone; chapter 29 distinguishes this estimate from its update. How many emissions and how much cost it will bring depends on the energy system: chapter 29 does the sums.

Words Must Bear the Weight of the Numbers

This chapter's conclusion can be put in one sentence: by 2036 the model produces no general crisis. But that sentence alone would be false. It has to be said in full: no general crisis, and at the same time marked redistribution and a long adjustment in work. More than four extra points of unemployment in the advanced economies for several years, losses concentrated on the poorer half, a gap still open at the end of the horizon. Separating the first half of the sentence from the second would mean telling a different book.

For the poorer half taken on average, the 2036 loss is small: comparable to a few months of normal growth. For someone who loses their job, the individual fall can be much larger. For some,

cheaper services soften it; for others, lost income and experience never gained weigh far more.

The next chapters show how the rules of the calculation were corrected, and which assumptions drive the conclusions most.

Chapter 15

Revising the Model

In 1974, speaking to the graduates of the California Institute of Technology, the physicist Richard Feynman said that the first principle of science is not to fool yourself, and that you are the easiest person to fool. This chapter tells how this book's model, for a while, fooled the person building it, and how the deception was discovered.

A chart can change because the data change, or because a relationship in the model is corrected. Before comparing two versions, you need to know which of the two has happened. Here we look at the rules that made labor demand recover too easily, and at what changed when they were corrected.

The Versions of the Model

Development went through three phases. First an illustrative exercise, with simple assumptions and some errors. Then a first structured version, in which the recovery of employment was partly imposed by the very form of the equations. Finally revision v2, the one used in the book. The age module is a later extension, separate from these versions.

The comparison below puts the illustrative exercise and the structured version side by side. Watch the units: the "median household" of the old exercise and the average of the poorer half in the new model are not the same thing.

Table 32 – From the illustrative model to the structured version. The old −6% concerns the median household in the central case; the new −1.0% concerns the bottom half and the median of the draws. GDP percentages are differences relative to the world without AI in 2036.

Feature	Earlier illustrative model	Structured model
Households in 2036	Median household −6%, central case	Bottom half −1.0%; 92% of draws negative
World GDP 2036	+9%	+1.3% (0.3–3.5%)
Productivity	Arbitrary gain per automated unit	Savings on automated tasks
Automation	Up to 38% of hours eliminated directly	Exposure × substitutable fraction × adoption
Capital income	Concentration chosen ad hoc	WID shares, evolving with saving
Investment	Absent	Responds to the profit share
Parameters	Chosen without calibration	Calibrated or justified ranges; 4,000 draws
Accounting	Not checked	Balances to numerical precision
Labor market	Employment without limit, later corrected	Employment limited by supply; wages react to shortages

The most important change concerns productivity. The initial exercise gave a large output gain for each automated unit, chosen without any basis. The later version links it to the cost saving on tasks actually automated, which is only part of their value. The extra output to distribute shrinks, and with it the size of the effects on households. The direction, however, does not change: gains concentrate at the top and the poorer half loses. The old gains and

old damage turn out to be overstated by a factor of five to ten; the table shows which quantities this judgment rests on.

A correction can reduce one risk and make another more visible.

A Suspicion Becomes a Check

During development an explicit check was made for possible biases in favor of the technology, in the assumptions, the equations and the interpretations. It was not an independent peer review. But it started from the right question: is the model saying what the numbers show, or what its builder would prefer to hear? The check found six problems.

The first concerned the recovery of households by 2046. It depended on the assumption that labor demand would grow faster than supply. With neutral demand growth, the recovery disappeared. A conclusion presented as a result actually came from a very favorable initial choice.

The second concerned a word: "temporary". In the discontinuity scenario households were still worse off in 2046. Calling that damage temporary was not consistent with the numbers.

The third was circular reasoning. The link between weak demand and profits had been judged unimportant on the basis of a parameter chosen in the previous version: the model was confirming what it had been told. The new model evaluates it with Sobol sensitivity analysis, over a range justified by the literature. Chapter 16 shows the results.

The fourth concerned the scoring of scenarios. Inclusive growth had been considered more plausible than discontinuity, even though it required policies not yet adopted. The comparison was redone giving the two scenarios equal plausibility. A desirable

measure cannot be treated as already in place just because it improves the ending.

The fifth was an asymmetry: unmeasured benefits, such as cheaper services, were emphasized, while unmeasured harms were neglected. Since then the book has kept the welfare adjustment separate and discusses both sides, including job quality, costs and material constraints.

The sixth did not concern an equation but words. Expressions such as "a large minority" or "slightly worse" softened what the model's own numbers showed. Revising the language serves to respect the results, neither inflating nor softening them.

During construction two further errors were corrected. The first: exposure to AI was held at today's level, excluding by construction any growth in capabilities. The parameter that widens exposure was introduced, up to 60 percent more over the period; and it is this limit that still excludes scenarios of almost total automation. The second: a calibration error hid the cost of support policies. In the corrected version the interventions start in 2027, with the same 2026 calibration, and chapter 13 uses these results, in which the cost of support and the extra debt are visible.

Jobs That Came Back on Paper

Then there was the most important error: the rule that calculated labor demand.

An example with simple numbers. With a wage bill of a hundred and a wage of ten, you can pay ten people; with a wage of five, twenty. That is arithmetic. But arithmetic does not say which activities need those twenty workers.

In the first structured version, labor demand was obtained exactly that way: wage bill divided by wage. If the wage fell, more people could be paid, and the model hired them. The mechanism

brought labor back even into tasks that AI now did more cheaply. Demand returned towards supply by construction: the model, in other words, had already written the happy ending into its rules.

The revision instead derives the labor required from technology and output: how many tasks remain for people per unit produced, valued at the wage of the world without AI. Then wages influence this demand with an elasticity, at constant output, between 0.15 and 0.75. A lower wage brings back only part of the lost work. The equations are in chapter 4; here is the reason for their form. The old relationship gave wages enormous power over jobs; the new one allows a partial recovery, linked to the human labor that production still requires.

There is also a cross-check: with an elasticity equal to one and no profit sharing, the new structure reproduces the old one to the computer's precision. And in the revision, profit sharing raises pay without touching the mechanism that determines jobs.

The World Without AI Can Be Wrong Too

Two checks must be kept apart. The first concerns how much wages react to unemployment: if they react little, the cost of labor recovers few jobs. The second concerns the comparison world: if it deteriorates because of a construction error, AI looks better than it is.

The reaction of wages to unemployment is measured by the slope of the wage Phillips curve. The old range ran from 0.2 to 0.8; the revision uses 0.05 to 0.5. The anchors are an American estimate that falls from about 0.45 between 1991 and 2008 to 0.14 between 2009 and 2015, not significantly different from zero, and the flattening of the curve after 2008 in ten advanced economies. After the great crisis, in short, wages stopped reacting much to unemployment.

And an error in the world without AI was corrected. There, wages followed productivity while the labor share kept falling for reasons unrelated to AI: the wage grew faster than the wage bill, and unemployment rose even in the comparison world. In the first edition it reached about 6 percent worldwide in 2046 and 7.8 in the advanced economies with very rigid wages. In the correction, the reference wage also follows the erosion unrelated to AI, and the world without AI stays at structural unemployment whatever the reaction of wages.

With the old ranges, this correction moves the medians by no more than 0.2 points. But it removes the small negative effects of AI on unemployment after 2040, which arose only from the comparison with a world without AI that was deteriorating by itself. A favorable effect of AI could be produced by a wrong term of comparison. That is why both trajectories must always be checked: shifting the reference by mistake can make a recovery appear that AI did not produce.

What Changed in the Results

Three other ranges were revised. The delay of new tasks went from 2–8 to 4–10 years. New tasks per task displaced from 0.25–1 to 0.25–0.85. The automated share from 0.15–0.7 to 0.25–0.7. The choices are justified by historical delays and by recent experience of activity creation.

Table 33 – Results of the first edition and of revision v2. The effects are differences relative to the respective no-AI references; pp denotes percentage points. The 2036 ranges are 5th–95th percentiles; for unemployment in 2046 only the 95th is reported in addition to the median. The scenario frequencies should not be added up as necessarily exclusive categories.

Result	First edition	Revision v2
ADV unemployment 2036	+1.2 pp (−0.1; +3.8)	+4.3 pp (+1.4; +10.7)
ADV unemployment 2046	−0.2 pp; 95th +0.7	+1.4 pp; 95th +9.0
Median of individual peaks	+1.8 pp	+6.3 pp
Frequencies A / B / C	10% / 70% / 19%	8% / 23% / 74%
Bottom half 2036	−0.6%; 90% worse	−1.0%; 92% worse
Bottom half 2046	−0.4%; 59% worse	−3.3%; 89% worse
World labor share 2046	−2.6 pp	−5.9 pp
ADV debt / GDP 2046	+2.6 pp	+11.2 pp
Package, bottom half 2036	+1.8%	+3.0%

The difference is stark. In the first edition, additional unemployment in the advanced economies was 1.2 points in 2036 and became slightly negative in 2046, minus 0.2: a problem that solved itself. In the revision it is 4.3 points in 2036 and remains at plus 1.4 in 2046, with one future in twenty reaching plus 9. For the poorer half, the 2046 loss goes from minus 0.4 percent, in just over half the simulations, to minus 3.3, in almost nine out of ten.

The frequencies of scenarios A, B and C in the revision add up to more than 100. The book keeps the calculated values and warns that they are not exclusive categories. A requires that the poorer half not be worse off; B combines a loss of less than 2 percent with an unemployment increase of less than three points; C applies with a loss of at least 2 percent or at least three points of extra un-

employment. A and C can therefore overlap, when consumption holds up but unemployment exceeds the threshold. The scenarios have a chapter of their own, chapter 21.

The chart isolates two medians from the table, without reproducing their ranges. It shows the shift from a slight final recovery to additional unemployment that persists.

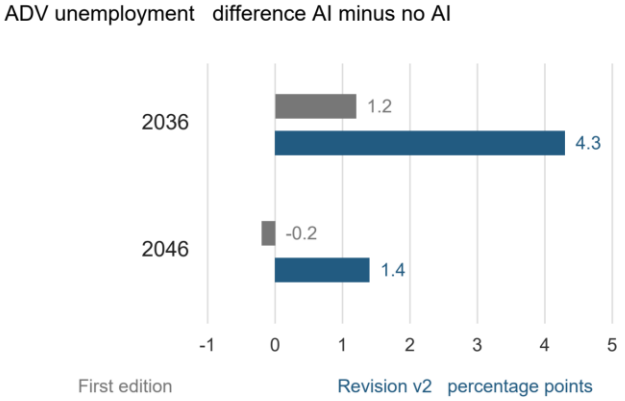


Figure 23 – Additional unemployment in the advanced economies in 2036 and 2046. Comparison of medians between the first edition and revision v2, in percentage points. It is a teaching redrawing of Table 33, not a new simulation and not a comparison between observed data.

Most of the change comes from rebuilding labor demand; the new wage range adds to it. The first edition's conclusion, that the problem would be solved in the early 2040s, no longer holds. The distributional losses keep their direction and become larger: for the poorer half the final gap grows instead of narrowing.

A more worrying result is not truer just because it is worrying: it is checked like the others, and chapter 18 says with which data.

The next chapter looks at which parameters explain most of the variability in the results, and where new evidence could change the picture.

Chapter 16

Sensitivity Analysis

A loaf of bread turns out badly. Is it the yeast, the flour, the oven or the rising time? If you change one thing at a time, you learn something. But not everything: a little yeast and a long rise can give the same bread as a lot of yeast and a short rise. Sometimes the result depends not on one ingredient, but on how the ingredients meet.

This book's model has twenty-one uncertain ingredients. When a result changes together with a parameter, the useful question is how much it changes, and in combination with which other assumptions. Global sensitivity analysis studies this. It serves to focus checks on the mechanisms that weigh most on the results.

One Number at a Time Is Not Enough

You can hold everything fixed and change a single parameter: bring adoption forward, increase the share of replaceable tasks, lengthen the time needed for retraining. The comparison gives a local answer, around the chosen combination. But the same test can give a different result when the other conditions change too. Early adoption has different effects if new tasks arrive at once or after many years. The same fall in wages brings back little work if firms react little to the cost of staff, and a lot if they react strongly. Two parameters interact when the effect of the first depends on the value of the second. That is why a series of isolated tests is not enough.

The method used bears the name of the Russian mathematician Ilya Sobol, who proposed it in the early 1990s. You look at how much the results scatter around their mean, that is, the variance, and try to attribute that scatter to the assumptions that changed. The results here are differences between the two worlds: consumption, GDP, unemployment, debt-to-GDP ratio.

The first-order index measures the contribution of a parameter on its own. The total index also includes the interactions in which that parameter takes part, and it is the one shown in the chart. A value of 0.48 means that the variability linked to that parameter, interactions included, explains about 48 percent of the variance of the result.

And the index does not say in which direction a parameter pushes. A parameter can matter a great deal in separating good and bad futures without the index saying which it makes good: that requires the economic mechanism and other comparisons. Nor do total indices add up like slices of a pie: an interaction can be attributed to several parameters at once.

The 4,000 main draws describe how the results are distributed. To separate the contributions of the parameters, a purpose-built set of comparisons is needed instead. The technical detail is this: a Saltelli sampling design, with a base sample size N of 1,024 and twenty-one parameters. The code calculates first-order and total indices, not all second-order ones. The construction requires N multiplied by the number of parameters plus two: 1,024 times 23, that is, 23,552 combinations. Each is calculated in the two worlds: 23,552 pairs, 47,104 paths in all.

The design takes its name from Andrea Saltelli, an Italian statistician who worked for many years at the European Commission's Joint Research Centre in Ispra, on Lake Maggiore, and who spent much of his career repeating a warning: a model must be questioned by changing all the assumptions together, not one at a time, otherwise a result that depends on a lucky combination is mistaken for a robust one.

The two analyses answer different questions: sensitivity separates contributions to variance, the main Monte Carlo explores the distribution of results.

What Decides the Results

The chart should be read one column at a time. Darker colors indicate higher indices. Each column is a result in a given year: consumption of the poorer half in 2036 and 2046, GDP in 2036, unemployment in the advanced economies in 2036, and their debt in 2046. Rows are compared holding result and year fixed.

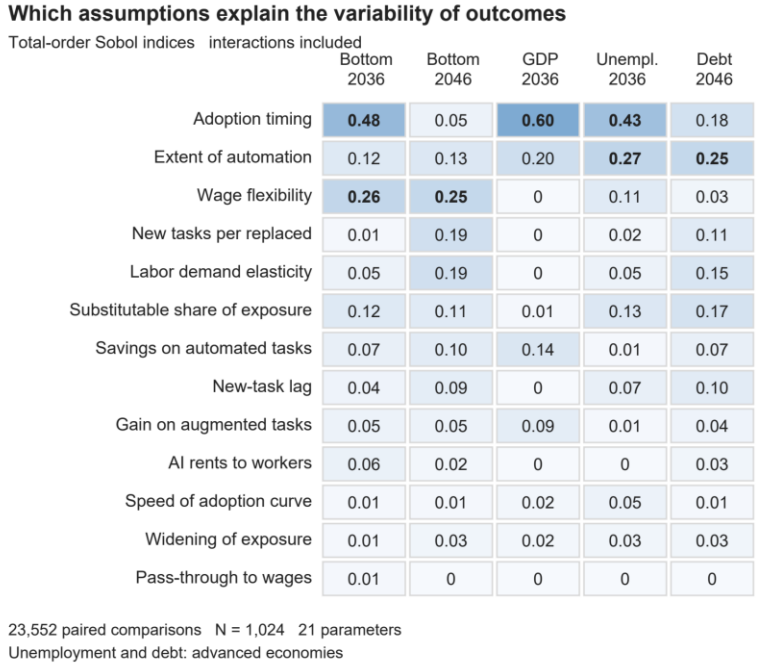


Figure 24 – Global sensitivity. Thirteen parameters, five outcomes and all printed values are shown. These are total-order indices, including interactions. Bottom half denotes the group's consumption; unemployment and debt concern the advanced economies. Of the twenty-one inputs, thirteen are shown. The two aggregate-demand parameters, excluded from the matrix, each explain at most about 0.015 of the variance. The sobol.json file also keeps unrounded estimates and the numerical uncertainty of the indices.

In 2036 a single parameter dominates: the year in which adoption accelerates most. Its total index is close to 0.48 for consump-

tion of the poorer half, 0.60 for GDP, 0.43 for unemployment in the advanced economies. A large part of the scatter at that horizon depends on where 2036 falls along the diffusion curve: at the start, in the middle, or when the race is already over. A process under way and one almost complete look different on the same date. That is why the book also looks at 2046.

In 2046 the timetable matters much less for the consumption of the poorest: the index falls to about 0.05. What weighs more is wage flexibility, about 0.25, new tasks for each task displaced, about 0.19, and the response of labor demand to wages, also about 0.19. In the long run, in other words, it no longer matters when the technology arrives, but what happens to work afterwards: whether new tasks appear and how wages react.

And here there is a delicate point. The appearance of new tasks does not guarantee that the people displaced can perform them: a task can arise elsewhere, require other skills, become accessible after years. And it is precisely on this capacity that the historical evidence is weak. In the long run it is the assumption that weighs most, and the one we know least about: that is where research is needed.

Wage flexibility has indices of about 0.26 in 2036 and 0.25 in 2046 for consumption. And it is not a free remedy. Lower wages can recover some of the work, but the wage is also the income of those who work. An adjustment can reduce the gap between demand and supply and, at the same time, reduce how much a household receives for each hour of work. And the elasticity of labor demand says how much work comes back when the wage falls: if the response is weak, a wage cut brings back few jobs. Task creation and the response to the cost of staff are different mechanisms, and higher output guarantees neither.

For unemployment in the advanced economies in 2036, after the adoption timetable come the extent of automation, about 0.27,

and the part of exposure that can be substituted, about 0.13. Wage flexibility weighs about 0.11. These are indices for the same result and should not be confused with those for consumption in 2046.

The distinction between exposure and substitution remains essential. An exposed task can be helped by AI without removing human work. The substitutable share says how much of that exposure becomes replacement; the extent of automation and the adoption path decide how much technical capability actually enters production. They are different steps in the same chain.

In the 2046 debt column, the extent of automation has the highest index among those shown, about 0.25. Shifting income from labor to capital changes the tax bases and the benefits needed during the transition. The fiscal result does not depend only on extra GDP, as chapter 11 showed.

The other rows retain information that a four-number summary would lose: savings on automated tasks, the delay of new tasks, the gain on assisted tasks, the share of AI rents reaching workers, the speed of the adoption curve, the widening of exposure, the pass-through of gains to wages.

The Loop That Weighed Little

There is a word in economics that means two different things: demand. Demand for goods and services is spending on products, and the response of output when that spending falls short. Labor demand is the work firms require; its response to wages appears in the matrix and matters mainly for consumption in 2046.

The analysis checks two parameters of the goods-demand loop: the response of investment to profits and the pass-through of a demand shortfall to output. Each explains at most about 1.5 percent of the variance of any result. That is why the chart does not show them. The loop this research started from, lower wages, less

consumption, less output, further losses, exists in the model's structure. But with the chosen ranges it weighs little compared with the other mechanisms.

The book keeps this result too, even though it complicates the intuition it started from. Research loses value when it keeps only what confirms its initial fears. The risk to work can be persistent and, at the same time, have a main cause different from the one first imagined: not a collapse in demand, but the speed of displacement and the slowness of new tasks.

And a sensitivity ranking is not a political program. A decisive parameter can be very hard to govern; a less important one can be changed cheaply. The measures in chapter 13 were evaluated with paired comparisons, not with these indices.

The ranking does, however, say where to look for evidence: the adoption timetable, the activities that reappear, the real possibility of performing them. The next chapter adds what a parameter matrix cannot describe: the relationship between economic hardship, access to services and trust in institutions.

Chapter 17

Social Cohesion and Institutions

On 15 January 2021 the government of the Netherlands resigned en bloc. Not because of an economic crisis or a war, but because of an administrative system. For years the Dutch tax authority had accused tens of thousands of families receiving childcare benefits of fraud, demanding enormous sums back. Many were innocent. Among the criteria used to identify "high-risk" cases was dual nationality. Many families found themselves having to repay sums they could not afford. And for years, those who tried to object found no one with the power to correct the error.

It is the extreme version of something that can happen anywhere: a service can respond faster without solving more problems. When an administration automates its communications, you have to distinguish the activity performed, the ability to access it and the result. And the distinction also concerns trust: receiving an automatic confirmation is useful, but knowing who can correct a mistake is another part of the relationship with institutions.

An Answer Is Not a Solution

When an administration automates its contacts, the number of cases processed rises, and that is the statistic that will be shown. But it is not enough. You need outcomes, errors, appeals, accessibility, differences between users. Messages sent and problems solved are not the same thing: a system can improve both, or increase the first while leaving the second unchanged. The efficiency of the procedure and the effectiveness of the service must be looked at separately. And responsibility can be recognized only if there is a contact person, the power to intervene and a way to communicate the outcome.

Institutions stand between economic hardship and social tension, and trust can be studied. In 2024 the OECD published a survey of trust in public institutions in thirty countries, collected in

November 2023. 39 percent of people said they had high or moderately high trust in their national government; 44 percent low or no trust. Among those who felt they had a say in what the government does, 69 percent trusted it; among those who felt they did not, 22 percent. Source: OECD, OECD Survey on Drivers of Trust in Public Institutions, 2024 Results, https://doi.org/10.1787/9a20554b-en.

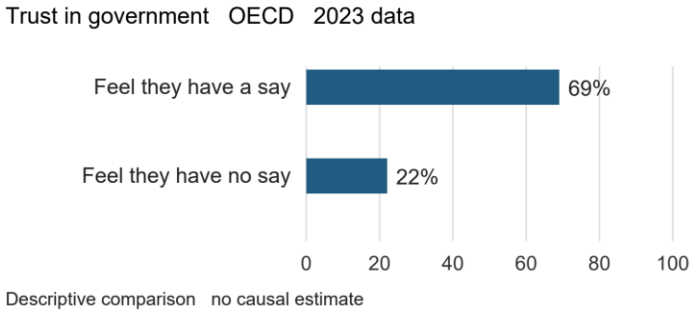


Figure 25 – Trust in the national government among people who report having or not having a say in what it does. OECD survey 2023, published 2024, thirty countries: 69% and 22%. Descriptive comparison between groups of respondents, not the causal effect of an intervention and not a measure of the probability of protest.

Sixty-nine against twenty-two: an enormous gap. The link runs both ways: those who feel heard trust more, and those who trust feel more heard. And to measure whether people really have a say, counting the meetings organized is not enough. A public meeting can be open to everyone and held when those who work cannot attend. A consultation can collect opinions without ever saying which ones counted. Access, understanding, response and influence on decisions all count. The figure widens the book's question: beyond protecting income and work, what matters is how institutions treat those going through the change.

Protest Is Not Violence

When AI and work are discussed, sooner or later someone evokes riots. It pays to be precise. Gathering to discuss, demand or protest does not show violent intent. Peaceful protest, unrest, crime and political violence have different causes and meanings. A march can exercise a right and make a difficulty visible; an assault is another problem. Putting everything into an index of "violence and social instability" makes democracy look pathological.

The economic model contains no verified formula that turns each point of unemployment into a certain number of protests, and the book does not propose one without a basis. Chapter 37 does something else: it declares a rule and follows its consequences, without presenting it as an estimate. What the model provides are conditions that can matter: losses concentrated in the poorer half, young people struggling to get in, a long unemployment peak, differences between regions. Between economic hardship and a collective response lie social networks, organizations, expectations, protections, the possibility of appeal. The ability to take part also depends on material conditions, such as times and transport.

What does research on unrest say? Hadzi-Vaskov, Pienknagura and Ricci, in a 2021 IMF working paper, analyzed quarterly data from 89 countries from 1990 to 2019, excluding states considered fragile. They use an unrest index built on newspaper coverage. When the index rises by one standard deviation, GDP remains on average about 0.2 points below its path after six quarters, a year and a half; for a major event, about one point. The effects vary with the motivations and with the strength of institutions. Source: The Macroeconomic Impact of Social Unrest, IMF Working Paper 21/135, https://www.imf.org/-/media/files/publications/wp/2021/english/wpiea2021135-print-pdf.pdf.

The standard deviation measures dispersion: here it gives a scale to the change in the index. And these are data from before generative AI: they say how much unrest costs, not how much AI will cause.

The next table brings together five links between economic hardship and security, each with what the evidence shows and what remains to be proved.

Table 34 – Economic hardship and security. Comparison of five steps. Historical and technical evidence is not an individual forecast of behavior; the legal reference is specified in the text with Article 5 of the European regulation.

Step	**Evidence in the model**	**What remains to be shown**
Economic hardship and unrest	Socioeconomic protests are followed by sharper GDP contractions than political ones. Strong institutions and fiscal space soften the damage (Hadzi-Vaskov and others, 2021).	That job losses due to AI cause unrest: the analysis identifies no events attributable to this channel.
Unrest and the economy	One-standard-deviation increase in the IMF index: GDP about 0.2 points below the reference after six quarters. Major event: about one point.	Long-run effects and transferability to the future of AI.
Budgets and the automation of security	Plausible incentive: security budgets face wage pressures similar to other services.	Savings greater than equipment, maintenance, cybersecurity, supervision and errors. Robust evidence on costs is lacking.

Step	Evidence in the model	What remains to be shown
Law and AI surveillance	The analysis refers to the EU ban on real-time remote biometric identification for law enforcement in public spaces, from February 2025, with authorized exceptions; it also refers to public social scoring.	That other jurisdictions follow the same limits. The legal summary is specified in the text.
Reliability of facial recognition	NISTIR 8280: in many algorithms, false positive rates differing by factors of 10 to over 100 between demographic groups; some of the best show no detectable differences.	That installed systems use the best algorithms and thresholds. The historical result does not describe every current system.

Protests born of economic hardship are followed by sharper falls in GDP than political ones, and strong institutions and sound public finances soften the damage. But none of these data concern AI. They say something that also holds for the future: the bill for a badly managed transition is not paid only by those who lose their jobs.

And a small average loss can be an enormous loss in a specific place. The model works with large regions and does not see neighborhoods; two cities with the same unemployment can have very different transport, housing and services. But it flags three conditions. Losses are concentrated and can remain almost invisible in aggregate data. The unemployment peak towards the end of the 2030s coincides with the pressure of pensions and debt. And profits can move from the emerging to the advanced economies while work is lost in both. These conditions generate different demands and divergences between territories.

The Camera and the Error That Becomes a Queue

There is a risk less visible than protest, and perhaps more likely. A state with less revenue and more needs may be tempted to automate security: more cameras, more algorithms, fewer people. The danger is not a dramatic decision. It is a slide: more surveillance and fewer services, one step at a time, without anyone ever having decided it explicitly.

But a camera does not, by itself, mean facial recognition, behavior classification or automated decisions: what counts is how it is used, for what, and under whose responsibility. The book distinguishes four things. Helping an operator search an archive. Planning patrol shifts. Running patrol robots. Deciding a coercive intervention automatically. They are not equivalent. The first helps a person; the last introduces responsibilities and risks of a different order.

The saving, moreover, is not calculated by counting the staff replaced. It requires purchase, maintenance, updates, cybersecurity, human oversight, error handling, the lifetime of the systems. And the time freed can be used elsewhere instead of reducing spending. The economic comparison must include the cost of a reliable service and the necessary safeguards; and an error rate must be read together with the type of error, the number of checks and the consequences of an alert.

The American standards institute, NIST, distinguishes false positives, which mistake one person for another, from false negatives, which fail to recognize a true match. It also distinguishes one-to-one verification, as when unlocking a phone, from one-to-many search, as when comparing against an archive. In 2019, in report NISTIR 8280, it documented that in many algorithms false positive rates varied between demographic groups by factors of ten to more than a hundred, while some of the best showed no detecta-

ble differences. Performance depends on the algorithm and on image quality. Sources: NIST, Face Recognition Technology Evaluation, Demographic Effects in Face Recognition, https://pages.nist.gov/frvt/html/frvt_demographics.html; historical report NISTIR 8280, 2019, https://doi.org/10.6028/NIST.IR.8280.

Why does even a small error weigh? An example, unrelated to facial recognition. Ten thousand checks, in which only ten cases are really the ones sought. A system finds nine of them: a sensitivity of 90 percent, which looks excellent. But among the 9,990 negative cases it flags ten by mistake. Result: nineteen alerts, of which only nine are right.

$$Correct\ share\ of\ flags = \frac{True\ positives}{True\ positives\ +\ False\ positives} \times 100$$

Nine divided by nineteen: about 47.4 percent. More than half the alerts concern people who have nothing to do with it, with a system that "works 90 percent of the time". When what is sought is rare, even a few false alarms weigh heavily on the work of checking, and on the people flagged. Judging a system requires the population, the threshold, error rates and operational consequences: an isolated percentage is not enough. And an alert to be checked must never become, on its own, proof of guilt. In the Netherlands it had become exactly that.

The European Union has already drawn some boundaries. The regulation on artificial intelligence, 2024/1689, in Article 5 bans real-time remote biometric identification in publicly accessible spaces for law enforcement purposes, except for narrowly defined exceptions and under precise conditions. The bans have applied since 2 February 2025. It is not a ban on every camera or every biometric use. And the ban on social scoring systems does not

concern only public authorities: it covers the practices, described in the regulation, that lead to unfavorable treatment out of context or that is unjustified or disproportionate. Source: https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng, Article 5; European Commission, First rules of the Artificial Intelligence Act are now applicable, 3 February 2025. Other countries do not draw the same boundaries.

The book's conclusion is this. Assistance in which an identifiable person decides is governable. Coercion decided automatically is a line not to be crossed. But adding a human signature does not make any procedure reliable: whoever decides must be able to understand, check and change the machine's proposal, and have the time to do so. A purely formal check leaves the error intact. And whoever suffers an error must be able to reach someone with the power to correct it.

The rule that ties this chapter together is simple. Every automated decision that affects a person must leave a trace: the request, the person responsible, the timing, the way the outcome is communicated. Recording a request with a name and a date makes the step checkable. But then you have to see whether the commitment is kept, whether people really manage to access the service, and whether those left out appear in the data. And when several offices contribute to the same result, their dependencies must be kept visible: solving one piece can leave the obstacle intact. Without a trace, efficiency becomes powerlessness for those who suffer the error.

The next chapter tackles the most honest question a book like this can ask itself: what do we not know, and how could we find out?

Chapter 18

Model Limits and Empirical Testing

In April 1992 New Jersey raised its minimum wage from 4.25 to 5.05 dollars an hour. Neighboring Pennsylvania did not. Two economists, David Card and Alan Krueger, had a simple idea: survey 410 fast-food restaurants in the two states, before and after the increase, and count the employees. Textbook theory said that where labor cost more, fewer people would be hired. The data did not show that. The study, published in 1994, sparked a controversy that lasted thirty years, and in 2021 it contributed to the Nobel Prize in economics awarded to Card. Not for the result itself, but for the method: comparing two similar groups, one touched by a change and one not, to see what would have happened without it.

A book that makes forecasts must say where it could be wrong, and how that could be discovered. To know whether an intervention works, you have to establish beforehand what it should change: attending a course, learning, finding work, earning more are different outcomes. This chapter separates the model's limits from the observations that would be needed to test them, and builds the comparison with what would have happened otherwise.

Ten Things the Model Does Not See

The model's documentation contains ten warnings. Having come this far, they can be read as questions addressed to a real city. They delimit what the model's numbers allow us to say.

The first warning concerns the way machines and people substitute for each other, and the way people enter work. In the model, the elasticity of substitution between tasks equals one. The elasticity of substitution says how much the choice between machines and people changes when their relative price changes: a cheaper machine does not make all tasks substitutable in the same way. With an elasticity of one, when wages follow the world without AI, the labor share mirrors the labor content of tasks. If the elas-

ticity were higher, the labor share could fall faster, as the line of research by Karabarbounis and Neiman suggests; if it were lower, it could fall more slowly, as the estimates for manufacturing by Oberfield and Raval indicate.

The same warning concerns participation. Labor supply follows United Nations demography: it does not respond to the length of studies, to participation, to migration. And the model does not see the obstacles of skills, age and place that separate a displaced person from an available new task. This can make unemployment look shorter than it would be. Conversely, the lack of a participation response can leave among the unemployed difficulties that in reality would become inactivity. The age module adds entries, returns and discouragement, with parameters set by the book.

The second warning is the absence of distribution within groups. The average for workers can improve while displaced young people lose a great deal: the model does not reconstruct these differences and may underestimate hardship. The third is the absence of ages in the main model: data on youth hiring do not enter its equations.

The fourth concerns adoption: its curve is set from outside, and it does not change with wages, profits and rules. The fifth concerns the missing macroeconomic links: there is no inflation, exchange rates or trade, the interest rate is given, and offsetting a demand shortfall depends on an assumed policy response. The regions are linked only through AI profits. A city that trades, imports energy and borrows does not live within those links alone.

The sixth is geography that is too coarse: three regions, with China and India in the same group. It hides enormous differences in institutions, population, specialization and income.

The seventh, the most important, is the technological frontier. The model excludes an extreme transformative AI, able to automate most cognitive and physical activities within the horizon,

such as the one discussed by Korinek and Suh. Exposure can grow by at most 60 percent. It means that all the results in this book hold for a cautious scenario. If AI goes further, things will be worse.

The eighth concerns probabilities: the uniform parameter ranges are reasoned choices, not distributions estimated from data, and the frequency of an outcome in the simulations depends on those choices. The ninth concerns unmeasured well-being: cheaper services, health, energy and job quality can change lives without properly entering the main indicators. The tenth concerns calibration: many regional assumptions lie within the ranges of the literature, but that does not make them observed data.

Ten warnings of different kinds: some require new equations, others new data, others only more caution in reading the curves. Some make the model more pessimistic than reality, others more optimistic. The seventh, the weightiest, points in one direction only: in its cautious version, the model risks underestimating the problem.

And there is one more limit, which concerns words. The model counts who works, but it does not know what that work is like. Being employed does not tell you how predictable the hours are, how continuous the income is, what the job is like. The OECD framework on job quality distinguishes three dimensions: earnings quality, labor market security, and the quality of the working environment. On its own it does not measure the effect of AI on a person, but it says what would need to be measured. Source: OECD, Job quality, https://www.oecd.org/en/topics/policy-issues/job-quality.html.

An individual record should separately register real pay, desired and actual hours, continuity of income and working conditions; and, for the problem of access, transport and compatibility with training, defined before the research. Without adding them up into

an arbitrary index: better pay can coexist with unpredictable hours, and an average would erase exactly that conflict. To evaluate re-employment you would also need the duration of the new job and the distance from previous income. Being employed on the day of the interview does not prove the transition succeeded. Losing a job can leave scars: lower wages, fewer hours, giving up the search.

The Missing Comparison

Even without AI, customers, interest rates, the composition of the workforce and conditions in sectors change. Looking before and after describes a change; attributing it to a cause requires a credible comparison with what would have happened otherwise. It is the problem Card and Krueger solved with fast-food restaurants.

Take two groups of occupations: one more exposed to AI, one less.

Table 35 – Teaching example: employment rate, percentage of the group's population.

Group	Before (%)	After (%)
More exposed occupations	60	55
Less exposed occupations	70	68

The more exposed group loses five points; the less exposed loses two. The difference between the two changes is minus three points. Final levels are not compared, because the two groups started from different points: the changes are compared.

$$DID = (55 - 60) - (68 - 70) = -3\,pp$$

This is the logic of difference in differences. To give it a causal meaning, among other things the two groups must have moved in parallel in the absence of AI. Looking at earlier years helps reveal whether this assumption is contradicted, but it does not prove

what would have happened afterwards. If the exposed group also suffers a crisis in its sector, the difference will attribute to AI something that does not belong to it. And if AI also affects the comparison group, what the comparison measures changes.

And the release of a tool at the end of 2022 does not mean all firms adopted it that day. An analysis over time must distinguish potential exposure from actual use, look for anticipation effects, follow the composition of the groups and the uncertainty of the estimates. And one year's twenty-two-year-olds are not the next year's twenty-two-year-olds: following an age group and following the same people are two different questions.

Research Written Before the Result

The American signal on young people is the most important we have. But does it also hold in Europe and Italy, where work is more protected? The book proposes a protocol to find out. The 2019–2026 window serves to compare the years before AI with those after.

Table 36 – European and Italian protocol. Hypotheses to be tested, not results obtained.

Step	Protocol proposed in the model
Question	Since the end of 2022, has the employment of 22–25-year-olds in occupations exposed to AI fallen relative to that of less exposed peers in Europe and in Italy?
Data	Eurostat EU-LFS scientific microdata, through a research project; Istat continuous labour force survey; INPS data accessible to researchers for hiring and firm-level controls.
Exposure	ILO 2025 gradients linked to four-digit ISCO-08 occupations; detail and crosswalks to be checked in the authorized archives.

Step	Protocol proposed in the model
Method	Event study by age and exposure quintile, 2019–2026; difference-in-differences relative to less exposed occupations; firm × time fixed effects where the administrative data allow.
Outcomes	Employment, new hires, terminations and wages. Availability and definitions depend on the source.
Pre-specified hypotheses	1. Relative decline of young people in highly exposed occupations with little complementarity to AI. 2. No comparable decline above age 30. 3. Adjustment through hiring rather than terminations.
Expected contribution	Test whether the US signal is found in markets with stronger employment protection; provide direct evidence to calibrate the youth channel. Any claim to be first must be checked against the updated literature.

Writing down the hypotheses before looking at the results serves to resist a temptation: choosing afterwards the curve that proves you right. If the hypothesis on young people were not confirmed in Europe, that result should be published with the same clarity. Stronger employment protection changes the channel through which the market adjusts, not the fact that it adjusts.

Access to data is a concrete step. For Eurostat's scientific microdata, the organization must be recognized as a research entity, and the project must follow the required procedure. Istat and INPS have their own conditions, to be checked archive by archive. The code can be prepared now on the public documentation. This book has neither obtained nor analyzed those confidential data. Source: Eurostat, Access to microdata, https://ec.europa.eu/eurostat/web/microdata/access.

Valuable help would come from Istat's longitudinal labor force files, which follow the same people three months apart. They

make it possible to study moves from one state to another instead of confusing a balance with a story. They concern those who remain resident in the same municipality during the period, so those who move require care. They are a different possibility from the public Eurostat series used in the project, which has no ages. Source: Istat, Rilevazione sulle Forze di Lavoro, dati longitudinali a 3 mesi, https://www.istat.it/microdati/rilevazione-sulle-forze-di-lavoro-dati-longitudinali-a-3-mesi-di-distanza/.

First entry requires an even more careful definition. Not having a job in the previous quarter does not mean never having worked. You need a sufficiently long history, a reliable retrospective question, or an administrative archive with adequate coverage; and even a first recorded contract may have been preceded by unrecorded work. That is why category F in the module remains a classification of the model, not a verified first entry. And those who do not take part in a course must appear alongside those who complete it, with the reason when possible: otherwise you observe only those who managed to get in, and lose precisely the people who met the biggest obstacles.

Beyond young people, three gaps remain to be filled. The first is new tasks, the reinstatement that weighs so much on the simulated future and is so poorly known. The work on occupational classifications and new job titles, such as that of Autor and co-authors, should be extended to the age of AI. A new title is not enough: you need to understand which activities are really new, how many people they employ and whether they offset those that disappeared.

The second gap concerns taxes. There is no measure of how AI changes the composition of tax revenue: administrative data by type of income and sector would be needed. Seeing a sector's profits rise does not say how much the state collects: tax residence, taxable bases, deductions and the rules of the period all count.

The third gap is distribution within groups: linking the model to household data. Regional averages do not say who owns capital, how many in a family work, what support they receive, what costs they face. The same lost salary weighs differently in two different households. And the link requires a documented procedure and data protection: it is not obtained by giving everyone the average of their group.

In any case the evaluation must distinguish four questions: whether the opportunity exists, who manages to reach it, what they find once inside, and against which alternative the change is measured. These limits do not weaken the thesis: they say what would be needed to confirm or refute it. The next chapter returns to the sources, to see which studies concern those already working, which those trying to get in, and which only what the technology might do.

Chapter 19

Review of the Evidence

What do we really know today about artificial intelligence and work? Much less than the headlines say. But more than those who repeat that "nobody can know" are willing to admit. This chapter lines up the evidence. Not to add up percentages, as if they were all the same thing, but to understand what each one proves. Every study observes a certain population, in a certain period, a certain phenomenon.

The Chain and Its Links

At the start of this research there was a chain: less labor demand, lower incomes, weaker consumption, possible effects on profits, more need for protection; and then risks for trust and for the use of automated security. Every link seemed plausible. But a chain of plausible links is not a proven chain. Some links are supported by research without the whole chain ever having been observed. And shifting income towards those who spend less, for example, does not rule out investment or transfers offsetting part of the demand lost.

The book's model combines Acemoglu and Restrepo's task model with stock-flow consistent accounting, in the tradition of Godley and Lavoie. It carries uncertainty through 4,000 pairs of simulations, studies the sensitivity of results and compares nine policies on the same basis. Here we go back to the evidence that motivates it. The table brings together assessments made at different stages of development, with their bases.

Table 37 – Assumptions and assessments of the model. Assessments made at different stages of development, with their respective bases.

Question	Assessment and basis
Tasks and occupations	Supported: automating an activity can transform an occupation made up of several tasks. Transformation is more likely than complete replacement. ILO 2025; Acemoglu and Restrepo 2019.

Question	Assessment and basis
Cognitive work	Supported: AI extends automation beyond routine manual work. High exposure in administration, software, media and finance. ILO 2025; IMF 2024.
Substitution and new tasks	Plausible risk in the qualitative assessment and supported as a risk in the model. Historical reinstatement falls from 1.0 to 0.5 per unit of substitution after 1987 (Acemoglu and Restrepo 2019). The review describes broad exposure and an observed impact that is still narrow, with no observed aggregate deficit: youth entries are a first signal. In the model with rigid wages, ADV differential +4.3 points in 2036 and +1.4 in 2046. A collapse is overstated in the initial thesis; the transition is understated.
Distribution and demand	Plausible mechanism, unknown empirical size. The world labor share has fallen by 1.6 points since 2004; investment and transfers can offset lost consumption. In the model the deficit appears in few simulations and is contained.
Demand and profits	The collapse of profits is overstated and not supported within the parameter space explored. In the worst of the 4,000 simulated futures demand falls short by 2.5% of output; in 95 futures out of a hundred it stays below 0.8%, and in 7 out of ten it never falls short. The demand parameters explain at most 1.5% of the variance of outcomes. The concentration of profits is more robust than their collapse in the model (Chapters 5 and 16).
Taxation and welfare	Erosion of labor tax bases and greater need for protection: plausible and well formulated in the first assessment, understated in the conclusions. Debt already close to 100% of GDP, capital lightly taxed and mobile, AI users and rent recipients in different places. In the model the risk is manageable if taxes follow capital income (Chapter 11).

Question	Assessment and basis
Robotics	Partly supported. 4.66 million industrial robots in 2024; service robotics and humanoids still early-stage and expensive.
Security and institutions	The link between insecurity and surveillance is speculative in the qualitative assessment, plausible but not quantifiable in the model. Weak institutions can favor it; EU law limits it. The conclusions judge inadequate a reading of security as simple fiscal substitution: law and responsibility matter. Chapter 17.
Contribution-based welfare	Concern correctly formulated: pensions and insurance financed by contributions are the most exposed institutions according to the analysis.
First entry	Absent from the initial thesis; clearest observed signal according to the review of the evidence. Closing the entry door also threatens the formation of future skills.
Social stability	Appropriate caution: institutions mediate the effects, and a protest is not a collapse.

Two things emerge strongly. The first: the collapse of the economy was overstated. Demand holds up, profits do not collapse; on the contrary, they concentrate. The second: the transition was understated. The damage is not a collapse, it is a long adjustment in which those at the bottom lose and the young are left out. And first entry, which was not even in the initial thesis, has become the clearest signal.

A detail of accounting honesty. An earlier version of the summaries gave 1.9 percent of output as the maximum demand shortfall, another 2.5. Rerunning the 4,000 simulations of the current model, the largest shortfall, in the worst future and the worst year, is about 2.5 percent; in 95 futures out of a hundred it stays below 0.8, and in seven futures out of ten demand is never short. The 1.9

does not correspond to the current model and has been corrected. The conclusion does not change: no demand crisis.

The most important research on automation and work is that of Daron Acemoglu and Pascual Restrepo. They broke American production down into activities assigned to machines and activities assigned to people, and measured two forces: displacement, which moves tasks towards capital, and reinstatement, which creates new tasks in which people have an advantage. The second force exists. But it has no guaranteed speed.

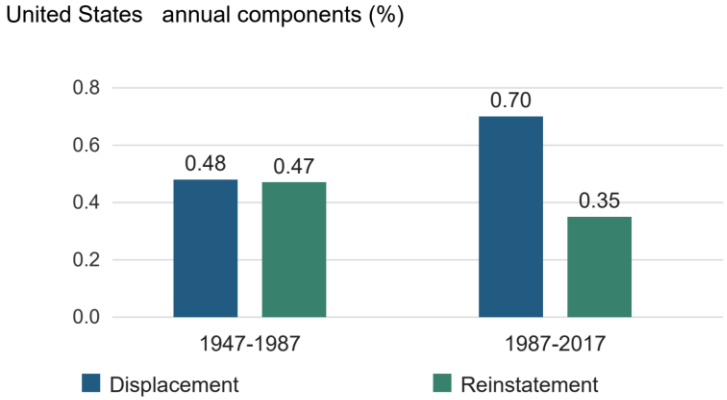


Figure 26 – Task substitution and reinstatement, United States. Annual components reported in the model from Acemoglu and Restrepo (2019), in percent; not job loss rates. Figure for reading the historical review. Source: Acemoglu and Restrepo (2019).

Between 1947 and 1987, in the United States, the two forces balanced almost perfectly: 0.48 percent a year of displacement, 0.47 of reinstatement. Between 1987 and 2017 displacement rose to 0.70, reinstatement fell to 0.35. For every task taken away, half a new one appeared, and the labor content of tasks fell by about 10 percent. In manufacturing, displacement reached 1.1 percent a year. These are components of a historical decomposition, not percentages of people fired each year.

In 2022 the same authors estimated that task displacement explains between 50 and 70 percent of the changes in the structure of American wages over forty years. It has already happened, before generative AI. And Korinek and Suh studied the race between automation and capital accumulation: in their simulations slow automation can come with rising wages, while fast automation can depress them long before automation is complete. The book uses their mechanism, but excludes the extreme transformative hypothesis from its simulations.

How Long a Machine Takes to Become an Organization

In 1987 the economist Robert Solow, who won the Nobel Prize that same year, wrote a sentence that became famous: you could see the computer age everywhere but in the productivity statistics. The computers were there; the measured gains were not. They came, but a decade later.

Being able to perform an operation is not the same as having an organization able to use it. Procedures, responsibilities and skills take longer than the tool. The economic historian Paul David showed this with electricity: it took about forty years for electricity to provide half the motive power of American factories. Comin and Hobijn, across 166 countries, found an average adoption lag of 45 years, shorter for recent technologies. Brynjolfsson, Rock and Syverson describe a productivity "J-curve": at first, firms invest in organization, skills and intangible assets that statistics measure poorly, and measured productivity understates the gain; later it can overstate it, when those investments start paying off without having been counted before.

And for AI the pace seems faster. An NBER study by Bick, Blandin and Deming reports that at the end of 2024 generative AI

assisted between 1 and 5 percent of hours worked in the United States, with reported time savings equal to 1.4 percent of total hours. It distinguishes how many people try a tool from how much of the work is actually assisted. Source: The Rapid Adoption of Generative AI, NBER WP 32966, https://www.nber.org/papers/w32966.

Four Percentages, Four Questions

The exposure measures in chapter 2 classify tasks and occupations with different perimeters. Eloundou and co-authors find that in the United States about 1.8 percent of jobs have more than half their tasks exposed to language models alone, and about 46 percent if the software accompanying them is included. Changing the technical perimeter changes the measure. And a job includes preparation, checking, assistance and handling the unexpected: automating part of it does not by itself decide the hours required by the rest, nor how many people are needed.

The table gathers ten measures, each with its source.

Table 38 – Measures and estimates from the sources cited. The rows have different populations, periods and methods; they do not form a single series or a uniform forecast. Sources: the publications indicated in the rows.

Measure	Value and source
Employment exposed to AI	About 40% worldwide, 60% advanced economies, 26% low income. IMF, Gen-AI and the Future of Work, 2024.
Occupations exposed to generative AI	One worker in four; 3.3% in the highest gradient. ILO Working Paper 140, 2025.
Employment gap for 22–25-year-olds in exposed jobs	19% below the path of less exposed peers; no comparable gap among experienced workers. With the yearly weights of chapter 3: −4.2% a year against −0.3%. Brynjolfsson, Chandar and Chen, August 2026. Descriptive.

Measure	Value and source
Productivity in customer support	+14% on average; +34% for novices. Brynjolfsson, Li and Raymond, NBER 31161. 2023 version, distinct from QJE 2025 in the text.
Total factor productivity from AI, ten years	No more than +0.66% cumulative. Acemoglu, Economic Policy 2025. Conditional estimate.
AI investment and US growth	About +0.5 points of GDP growth in 2025. IMF Annual Report 2026.
World labor income share	52.4% in 2024, from 53.0% in 2014. ILO WESO, May 2025.
Industrial robots in operation	4.66 million in 2024; annual growth +9%. IFR World Robotics 2025.
Job churn expected by employers to 2030	170 million created, 92 million displaced. WEF Future of Jobs 2025. Survey of expectations, not jobs already observed.
World public debt	Just under 94% of GDP in 2025; 100% by 2029. IMF Fiscal Monitor, April 2026. Level and projection distinct.

Each row answers a different question, about a different population, in a different period. Together they do not make a forecast. But they all point in the same direction: wide exposure, rapidly growing use, first signals among the young, a labor share already falling, more and more robots.

The customer support row uses the 2023 version of the study, discussed in chapter 8. The article published in 2025 in the Quarterly Journal of Economics updates the sample to 5,172 agents and the average gain to about 15 percent. Sources: NBER WP 31161, https://doi.org/10.3386/w31161; QJE, https://doi.org/10.1093/qje/qjae044.

And that study measures those already inside; the Stanford row, in its August 2026 version, measures the path of young people in exposed occupations. The two results coexist: more efficiency for those already working does not guarantee more access for those trying to get in.

From the desk to the country is a long step. Acemoglu estimated a total factor productivity gain of no more than 0.66 percent over ten years, combining how many activities are affected and how much is saved on each. It is not the 14 percent of the call center multiplied by all workers: total factor productivity concerns what is obtained from all inputs together, not the cases an agent resolves in an hour. Source: The Simple Macroeconomics of AI, NBER WP 32487 and Economic Policy 2025, https://doi.org/10.3386/w32487. Other estimates are much higher, and they can differ tenfold because exposure, diffusion, savings, complementarity and new tasks change. But all of them answer a question about output. Who will receive the income is a different question.

On physical work, Acemoglu and Restrepo estimated that in the American local markets studied, one more industrial robot per thousand workers reduced the employment-to-population ratio by about 0.2 points and wages by 0.42 percent. In Germany, Dauth and co-authors found that jobs lost in manufacturing were offset by jobs in services, but that the cost fell on young people who were supposed to enter, not on those already employed. The national total may not show the door closing in a sector. And the 4.66 million industrial robots in operation in 2024, up 9 percent, measure installed machines, not jobs lost.

Distribution Had Already Changed

The labor share was already falling before AI. Karabarbounis and Neiman document a decline of about five points in the global cor-

porate sector between 1975 and the early 2010s, and attribute about half of it to the ever lower price of equipment. The ILO series goes from 53 percent in 2014 to 52.4 in 2024, with a fall of 1.6 points since 2004. The decline came before generative AI.

How easily machines and people substitute for each other also varies across estimates: 1.25 for Karabarbounis and Neiman, between 0.5 and 0.7 for American manufacturing according to Oberfield and Raval. Autor and Dorn link the polarization of jobs and wages between 1980 and 2005 to the ever lower cost of automating routine activities. And there is a reassuring fact, which must be stated: according to Autor and co-authors, about 60 percent of American employment in 2018 was in occupations that did not exist in 1940. New work really does appear. The question is whether it will appear fast enough, and for the right people. And a new job title is not enough to measure new work: you need activities, numbers of people, time to build skills.

How much do households spend when they receive one more euro? Annual estimates of the marginal propensity to consume range from 0.2 to 0.6; in the simulations of Carroll and co-authors, from 0.48 for the poorest fifth to 0.06 for the richest tenth. The average propensity, total spending divided by income, is something else: about 0.97 for the poorest fifth and 0.48 for the richest, according to Fisher and co-authors.

And wealth is highly concentrated. According to the World Inequality Report 2026, the richest ten percent receive 53 percent of pre-tax income and own about 75 percent of wealth; the poorer half receives 8 percent of income and owns 2 percent of wealth. These shares motivate the division of households in the model. AI arrives in a world already unbalanced, and pushes in the same direction.

The state has its timescales too. Public finance depends on how income is composed and on pressures that already exist, such as ageing.

And there is history. In Britain between 1800 and 1830, according to Robert Allen, output per worker grew by 0.63 percent a year while real wages stood still. They recovered after 1830, and only after 1860 did they begin to grow faster than output. Thirty years of stagnation are most of a person's working life. That is why the path counts, and not just the end point.

Some responses have already been tried. In Finland, in 2017 and 2018, a basic income had small effects on employment and improved well-being, trust and economic security. In Alaska the permanent dividend did not reduce total employment, and it increased part-time work. In Kenya cash transfers led to the creation of businesses without reducing hours worked.

Short-time work schemes can preserve jobs, but they do not guarantee that workers move towards growing firms. Giupponi and Landais find strong effects on the number of employed and negative spillovers on employment growth in other firms. In Germany, in April 2020, almost six million people were on short-time work. Preserving an employment relationship during a recession and building a new path after permanent displacement are different goals.

Sector-focused training programs, studied by Katz and co-authors, raised participants' earnings by between 11 and 40 percent, durably, mainly because they placed them in better-paid sectors. Counting completed courses is not enough: selection, support and employer demand matter. And collective agreements can write rules: in 2023 the American screenwriters' union, the Writers Guild of America, won agreement that AI-generated material is not considered literary material, that its use cannot be imposed on

writers and that companies must disclose it. It shows that rules can be written.

A Signature Does Not Replace Verification

Comparing the assessments makes visible things an output measure does not summarize: youth entry, the geography of income, the continuity of social protection. Important variables are still missing: demography, adjustment times, where profits are booked, prices and investment as a source of demand. The IMF estimates that AI investment added about half a point to American growth in 2025. And ageing and labor shortages can coexist with people who lack a compatible job: place, skills and timing must meet. That people over sixty-five will outnumber children towards the end of the 2070s, as the demographic projection cited foresees, does not make compensation automatic today.

The measurement problems are concrete too. Putting GDP, profitability and a composite index on the same axis does not make their meanings comparable. Protest, crime and political violence are not a single social outcome. A world average can hide opposite regional paths.

The revision of the model points to four further mechanisms to study: the interruption of apprenticeship, with effects on the skills of future experts even after ten or fifteen years; demographic substitution in ageing economies; cheaper services, to be assessed together with job quality, surveillance, concentration and energy; the geographical concentration of rents. Chapters 13 and 22 compare the tools and who must implement them.

The review covers studies up to October 2026 from NBER, the IMF, the ILO, the OECD, the BIS and scientific journals.

A prestigious name does not make a number true, and a suspicion about the tool does not make every number false: sources,

data and equations are checked. The next chapter returns to public money: where income arises, where it ends up, and who can tax it.

Chapter 20

Tax Bases and Public Finance

In October 2021, 136 countries and jurisdictions, meeting within the framework of the OECD and the G20, agreed on a principle that for decades had seemed impossible: large multinationals will have to pay at least 15 percent tax on their profits, wherever they book them. It had taken years of negotiations, tax haven scandals and the realization that digital companies could sell in one country, book profits in another and pay almost nothing in either.

That agreement speaks of something that artificial intelligence will make even more important. A service can be produced in one place, the income it generates can arrive in another, and the taxes can be paid in a third. While the unemployed stay in the first. An overall account in surplus does not prove that every territory has the resources it needs. This chapter looks at this separation, a mechanism that concerns everyone.

A Euro That Changes Route

Go back to the example in chapter 11. In an imaginary country the levy is 28 percent on a euro of wages and 19 on a euro of profit. When a euro moves from wages to profit, the state loses nine cents. Not because anyone evades: the income has only changed form.

The example holds everything else fixed to isolate the mechanism. In the simulated economy, by contrast, output grows, incomes change together, and there are consumption, transfers and benefits. The nine cents are therefore not the total cost of AI. They explain why a more productive economy can collect less precisely on a certain base, that of wages, which finances pensions and social protection.

And a fiscal proposal, to work, must pass three tests: finding the income, managing to collect it, sending it where it is needed. In a

model they fit in one line. In reality they involve different people, timescales and institutions.

The following table summarizes the public accounts by scenario, in the advanced economies. The labor share is a level in 2036; the other columns are differences from the world without AI. The primary balance excludes interest; debt accumulates the effects of the budget over time. They are all called "points", but the columns are not interchangeable.

Table 39 – Public resources by scenario. Medians conditional on the scenario definitions, except for the overall row. Percentage points shown as pp; debt in points of GDP. The revenue rules are assumptions for unidentified countries and do not describe national systems.

Scenario	Labor share 2036 %	Δ share pp	Δ balance 2036 % GDP	Δ debt 2046 pp
A inclusive	51.2	−2.9	−0.7	+5.4
B polarized	52.2	−1.9	−0.4	+6.5
C discontinuity	48.7	−5.5	−1.2	+13.3
All: median	49.7	−4.5	−1.0	+11.2 (P5–P95: +3.4 / +26.0)

In the inclusive scenario the poorer half does not lose purchasing power. But the state loses all the same: the balance worsens by 0.7 points of GDP and debt rises by 5.4 points. The name "inclusive" does not promise that everything improves: there too unemployment and debt rise. Protecting households and protecting the public accounts depend on different mechanisms. A benefit for households can cost the budget; a better balance can come from less protection. The columns must be read together.

In the discontinuity scenario the labor share falls by 5.5 points, the balance worsens by 1.2 points of GDP and debt rises by 13.3.

In the median of all futures, debt is 11.2 points higher in 2046, and in one future in twenty more than 26. The comparison between scenarios B and C shows how far the conditions of households and of the budget can move from the overall summary in chapter 11. But the scenario rows do not add up: each summarizes different trajectories, and the accounts are checked within each trajectory, not by adding medians of different sets. And all this comes on top of a debt that, because of ageing, was already growing.

A Percentage Needs a Base

Chapter 11 calculated that balancing the accounts would require a levy of 41 percent on the income alone that moves from labor to capital. The base is the additional income shifted towards capital, not all existing profits. Changing the base of a percentage is like changing the denominator of an indicator: the question changes, and so does the result.

And to move from an account to a law, three things are needed: knowing to whom the income should be attributed, managing to collect it, foreseeing how firms will react. The experiment in chapter 13 is something else again: it raises the levy on all capital by five points, a much broader base, and returns the revenue through transfers. They are two different proposals, not two versions of the same one.

Then there is the place where income arrives. Imagine two countries. In the first, firms use an artificial intelligence service and reduce staff. The second is home to the company that sells that service. The first loses wages and contributions without gaining the same base of profits; the second collects profits and taxes. And the first has to pay the benefits.

The model represents it: part of the profits generated in the emerging and poor economies is booked in the advanced ones.

And it also holds within a country, between one province and another, between a city and its surroundings. A global account able to pay all obligations does not guarantee that every territory can. The place that loses taxpayers may be the very one that has to pay for training and services; the one that collects the rents may have different needs. The global balance hides this separation, unless it is read region by region. And the level of a levy and the right to collect it answer different questions: a rule can set the rate without giving every territory a sufficient base. It is exactly the problem that the agreement on the global minimum began to tackle.

Five Routes, With Their Costs

There are at least five possible routes for financing the transition, and none is free.

The first changes how much income from capital contributes. It follows the shift of income, and it is the most direct. But capital moves, and one would need to observe its mobility, the bases taxed and the real possibility of applying the rule.

The second uses taxes on consumption. Consumption does not need to know whether a euro comes from a wage or a profit: it reaches income wherever it is spent. But it weighs more on households that spend almost everything they have. Transfers are therefore needed to compensate, otherwise the state is protected at the expense of the poorest. Increasing revenue is not enough to show that the poorer half is protected.

The third looks at rents, land, scarce assets. It is right in theory and slow in practice: identifying and valuing them takes time.

The fourth broadens ownership, with public funds or worker participation. Capital income thus reaches households that previously had only wages. It is the route the model rewards most for

the poorer half. But it needs initial capital, governance rules and a decision on who takes part. Building it in reality is not free.

The fifth changes the tax incentives that today favor investments that replace labor. To evaluate it, investment, hiring and productivity must be compared. And it does not follow that every automation should be slowed down.

And the famous "robot tax", calculated on the number of machines? It has a basic problem: what do you count? A piece of software, a remote service and a mechanical arm do not have a common, economically equivalent physical unit. You would end up taxing what is visible, not what produces value, with possible effects on productivity.

Before Choosing

All these routes start from incomes, spending, ownership and rules, not from the number of visible machines. Every comparison must follow both the resources and who receives them. And none works if you start building it when the unemployment peak has already arrived.

A policy can be tested under different conditions without knowing which will come true. But you need to understand how the scenarios used to test it are built. The next chapter opens them, one by one.

Chapter 21

Scenarios to 2036

In the 1970s the oil company Shell began using a new method to think about the future. Instead of a forecast, it prepared several different scenarios, each internally consistent, and asked its managers: if this happened, what would we do? When the oil shock came in 1973, Shell was one of the best-prepared companies, because one of its scenarios had imagined exactly a sharp rise in the price of crude. Scenarios are not for guessing. They are for not being caught by surprise.

Scenarios about the future, however, often have names chosen in advance, reassuring or terrifying: "the world of abundance", "the collapse". In this book the opposite happens. Inclusion, polarization and discontinuity are names given to groups of results, not three futures with a known probability. The simulations start from the same conditions in 2026 and run to 2046; the definitions help to read the differences within that period.

The Name Comes After the Result

In the first version of the model the scenarios were combinations of parameters chosen by hand. In the revision the combinations are drawn at random, and then the 2036 results are classified. The name depends on what happens to the poorer half and to unemployment in the advanced economies. It is not a label assigned beforehand, to say that the technology is good or bad.

There are three rules. The future is inclusive if the poorer half does not lose purchasing power compared with the world without AI. It is a discontinuity if the poorer half loses at least 2 percent, or if unemployment in the advanced economies rises by at least three points. Everything else is polarized.

Table 40 – Scenario definitions and frequencies. All comparisons are with the world without AI in 2036. The clarification on B makes the code's classification explicit.

Scenario and frequency	Definition in 2036	Mechanism in the model
A inclusive 8%	Bottom half not worse off than without AI	Savings on automated tasks 38%, against 30% in B; more profit sharing. Unemployment still about +4 pp.
B polarized 23%	Bottom half loses less than 2% and ADV unemployment rises by less than 3 pp; A and C excluded in the code	Later adoption, fastest around 2038; automation of 38% of exposed tasks.
C discontinuity 74%	Bottom half loses at least 2%, or ADV unemployment rises by at least 3 pp	Automation of 51% of exposed tasks; adoption around 2034; new tasks after about 7 years.

The inclusive future occurs in 8 cases out of a hundred: firms save more on automated tasks, 38 percent against 30 in the polarized scenario, and share profits more. The polarized future in 23: adoption comes later, with maximum speed around 2038, and automation reaches 38 percent of exposed tasks. Discontinuity in 74: automation reaches 51 percent of exposed tasks, adoption accelerates around 2034, and new tasks arrive after about seven years.

Eight, twenty-three and seventy-four make a hundred and five, because the categories are not entirely separate: a future can be inclusive for consumption and at the same time have high unemployment, and so be in both A and C. The code calls B everything that is neither A nor C, and the definition of B must be read that way. The overlap is a property of the criteria, as well as of rounding. And it says a lot: even in the scenario called inclusive, unemployment in the advanced economies rises by about four points in

2036. Protecting the consumption of the poorest does not mean including them in work.

With cautious assumptions, three futures out of four end in discontinuity.

Six Windows on the Same Experiment

Each panel of the following figure uses the world without AI as its reference. Zero means no difference, not the absence of output or unemployment. The lines are medians within each scenario, not the path of a single economy.

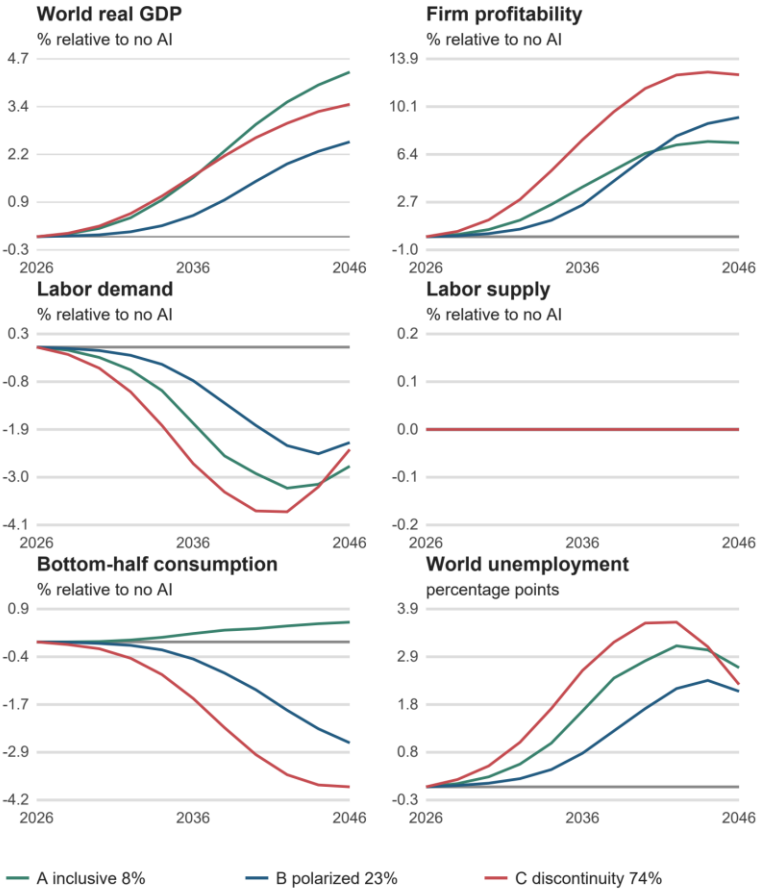


Figure 27 – Six indicators for A, B and C, from the model simulations. Medians of 4,000 paired simulations, years 2026–2046. Percentage differences except unemployment in percentage points; labor supply has a zero difference by construction.

Output grows in all three scenarios, while consumption of the poorer half falls in two out of three. More output does not automatically assign the extra income to the same people.

Profitability has a precise definition here: the capital share multiplied by the utilization of productive capacity. It is not the total profit of firms, nor a stock market price. In 2036 it rises by about 2.5 percent in the polarized scenario and by 7.6 in discontinuity, while consumption of the poorer half falls by about 0.5 and 1.5. The inclusive scenario distributes enough to make that consumption positive, even though labor demand remains below the world without AI. Profitability grows most precisely where poor households lose most.

Labor demand says how many people firms require at the current wage. Supply follows demographic projections and is identical in the two worlds. When they separate, unemployment arises in the model.

Table 41 – Definitions of the six indicators.

Indicator	Definition	Reading the result
World real GDP	Output at constant prices	Productivity of automated and assisted tasks; C automates more and produces more.
Firm profitability	Capital share × capacity utilization; not total profits	Rises when the labor share falls; highest in C.
Labor demand	Workers required at the current wage, compared with the labor force	Falls if substitution exceeds new tasks; recovers slowly, still below no-AI in 2046.
Labor supply	Labor force according to UN projections	Identical with and without AI.
Bottom-half purchasing power	Real consumption per capita of the bottom half of each region	Falls with the labor share, unless gains are shared; new tasks can support recovery.
Social pressure	World unemployment rate, in points	Measurable risk input, not a forecast of unrest. Composite index that could not be calibrated abandoned.

The label "social pressure" in the chart indicates unemployment. The old composite index of instability was abandoned because it could not be calibrated.

Two aggregate curves, moreover, can almost coincide while the labor chart shows a clear gap. To understand the difference you need to know the territory, the scale and the indicator.

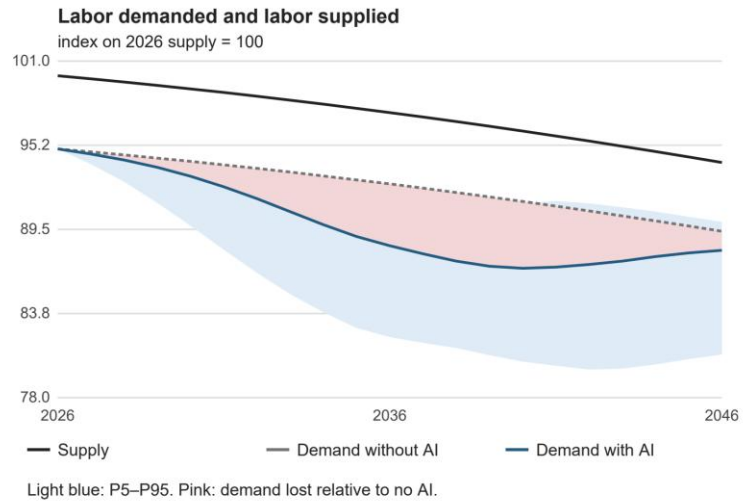


Figure 28 – Demand and supply in the advanced economies. 2026 supply = 100 for all series; median of 4,000 simulations. Band of demand with AI from the 5th to the 95th percentile; pink area between demand with and without AI. The indices are not percentage unemployment rates.

In the advanced economies supply falls slowly, and demand without AI already starts lower because of structural unemployment. AI opens an additional gap. In 2039 demand and supply are about 9.6 index points apart, of which about 4.7 are due to AI. These quantities are measured on 2026 supply; the unemployment rate for a year instead uses that year's supply. That is why the median peak of about 4.9 points of additional unemployment does not exactly coincide with the 4.7 index points.

In 2046 advanced labor demand remains about 1.3 index points below the world without AI. Part of the gap closes, but not by itself. And wages, in 2046, are 9 percent lower in the median and 25 percent lower in the worst future in twenty. Jobs can partly come back, and wages do not. Recovering employment and recovering income are two different tests.

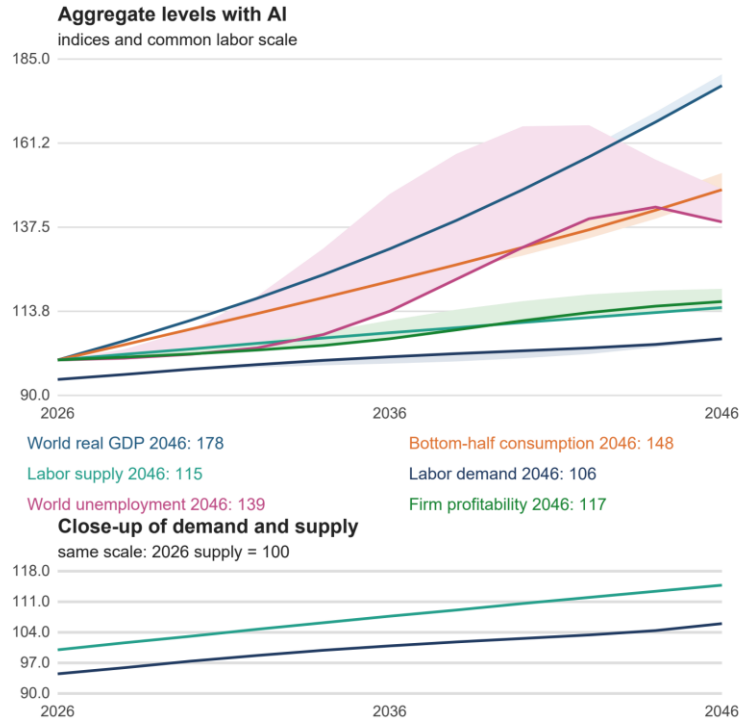


Figure 29 – World levels with AI: scenario B and the range between the A and C medians. Model simulations; below, an enlargement of the labor series only. GDP, consumption, profitability and unemployment are indices with their own base 2026 = 100; demand and supply both use 2026 labor supply = 100. The A–C range is not a percentile band.

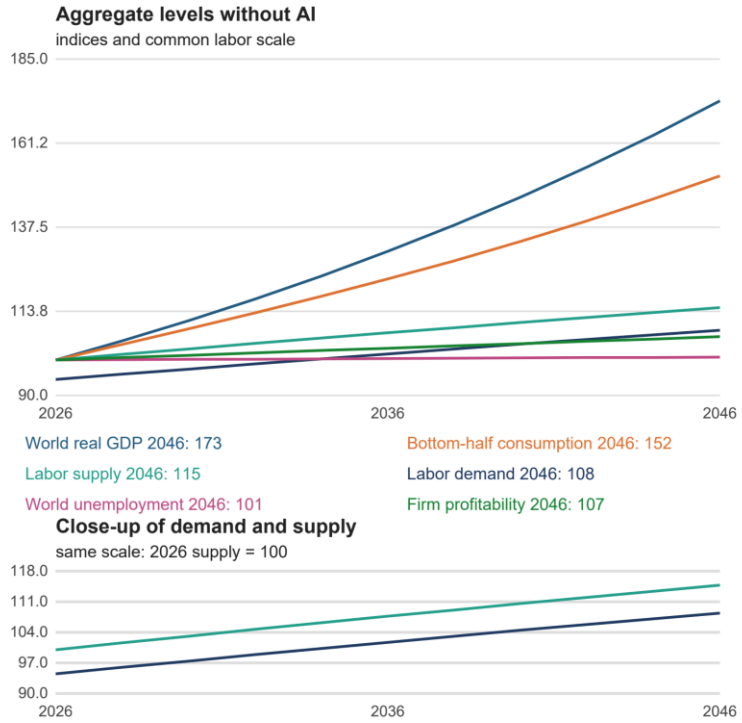


Figure 30 – The same indicators without the AI wave: median of the 4,000 simulations. Same base conventions as Figure 29 and enlargement of the labor series. Ordinary technical progress remains in the reference. The comparison between B with AI and the overall no-AI case is not the paired difference conditional on B: for the causal effect within the model, the difference figures are used.

The two figures use the same data; the second enlarges labor with consistent scales, so that the wider axis of the other indicators does not hide the gap. In 2046, with AI, world GDP reaches about 178, consumption of the poorer half about 148, profitability 117, supply 115, demand 106. Without AI: about 173 for GDP, 152 for consumption, 107 for profitability, 108 for demand, with supply still at 115.

The unemployment index close to 139 with AI does not mean 139 percent unemployed: it is the rate relative to its initial level, while without AI it stays almost flat around a hundred. The world gap exists even without AI, because of background unemployment, and with AI it widens. The scale that must also contain GDP close to 180 makes it hard to see, and the world as a whole softens the profile of the advanced economies, because the other regions adopt at different times. More output, less consumption for those at the bottom, more people without work: to see it you need to know territory, scenario, base and comparison.

A Snapshot Is Not the Whole Story

The classification uses 2036. But a snapshot contains neither the duration nor the order of events: to describe the transition you must follow the path before and after that date.

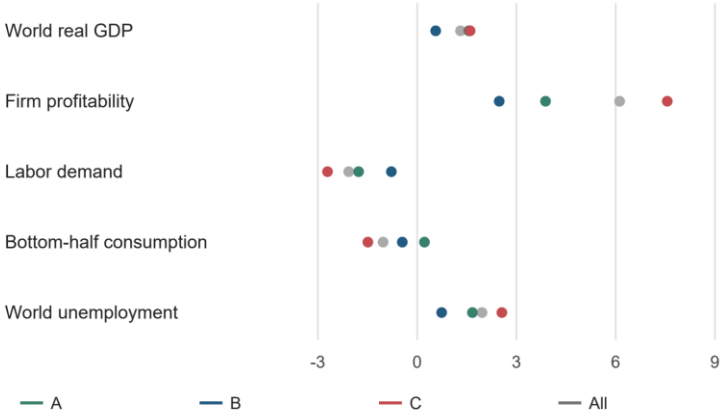


Figure 31 – World snapshot of 2036. Medians for A, B, C and the whole sample (gray). Percentage differences except world unemployment in percentage points. Colors as in Figure 27; different units made explicit.

Note a detail. World unemployment in the snapshot is not that of the advanced economies in the tables. Worldwide the increase is about 1.7 points in the inclusive scenario, 0.7 in the polarized one, 2.6 in discontinuity; in the advanced economies the corresponding values are about 4, 2.1 and 5.3. Calling both "unemployment" without saying where creates a false contradiction. Keeping the territory makes it possible to see where the pressure concentrates.

Table 42 – Scenarios, mechanisms and results. Group consumption in 2036 relative to no AI; unemployment in the advanced economies in percentage points. D and E are conditions within the main paths, not additional exclusive categories.

Scenario	Mechanism and frequency	Bottom half 2036 %	Top 10% %	ADV unempl. pp
C	Substitution exceeds new tasks for a decade; lower wages recover only part of the jobs. 74%.	−1.5	+3.8	+5.3
B	Later and more limited automation; output and profits grow, the labor share falls and entries weaken. 23%.	−0.5	+1.2	+2.1
A	More productivity and rents for workers; the bottom half gains, with higher unemployment. 8%.	+0.2	+2.6	+4.0
D demographic substitution	Condition in all scenarios: ADV supply −0.2 / −0.4% a year; automation covers shortages.	–	–	–
E stalled adoption	Reliability, energy and integration act as brakes. 12% of cases with GDP <+0.5% in 2036; bottom half −0.4%.	–	–	–

The discontinuity scenario has a clear logic: for a decade substitution runs faster than new tasks, and lower wages recover only part of the jobs. The poorer half loses 1.5 percent, the richest ten percent gains 3.8, unemployment in the advanced economies rises by 5.3 points. In the polarized scenario automation comes later and is more limited: output and profits grow, the labor share falls, entries weaken. In the inclusive one, productivity and rents go more to workers, and the poorer half gains, albeit with more unemployment.

Two conditions run through all the scenarios. The first is demographic substitution: in the advanced economies the labor force falls by between 0.2 and 0.4 percent a year, and some automation can cover real shortages. But covering a shortage does not guarantee that those displaced can perform the tasks still required. The second is adoption slowed by reliability, energy and integration costs: in 12 percent of the simulations the GDP gain in 2036 stays below 0.5 percent. Even there, however, the poorer half loses about 0.4.

The word "discontinuity" depends on a threshold. With four points of unemployment instead of three, its frequency falls to about 60 percent; with five, to 50. The unemployment criterion alone takes 46 percent of cases above three points. Labels depend on a threshold; continuous numbers do not: an increase of 2.9 points and one of 3.1 are almost the same for the economy, even if they fall on opposite sides of a threshold at three.

The People the Median Does Not Meet

The following figures use all the simulations, without limiting them to the labels of the three scenarios.

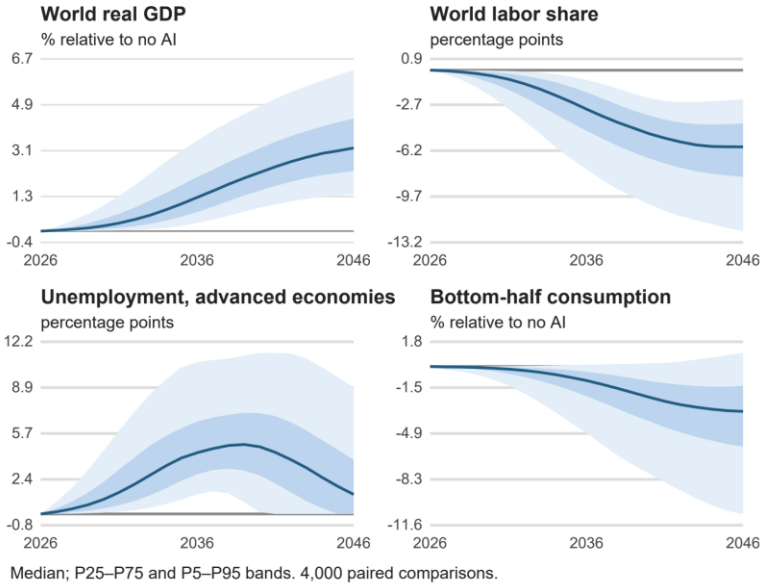


Figure 32 – Aggregate trajectories. Median and P25–P75 and P5–P95 bands of 4,000 paired comparisons. Percentages for GDP and consumption; percentage points for labor share and advanced-economy unemployment. The bands represent uncertainty over the parameters explored. Model simulations.

Output, labor share and consumption must be followed together: the improvement in the first curve does not cancel the others, and the bands show how much the paths scatter. GDP rises in almost all futures, the labor share falls in almost all, unemployment rises and then falls without returning to zero by 2046. And the regions' timescales are different: support must also be judged by how long it lasts, not only by how much it is worth in one year.

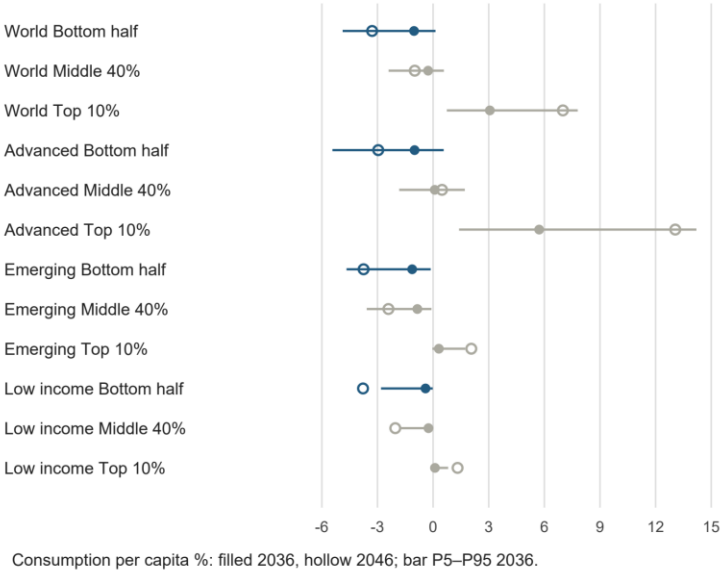


Figure 33 – Income groups and regions. Real consumption per capita relative to no AI. Filled dot: 2036 median; bar: 2036 P5–P95; hollow dot: 2046 median. The groups are defined within each region and do not follow individual people over time. Model simulations.

And here is the book's clearest result. In 2036 the poorer half loses in all regions, in the median. In the advanced economies it loses about 1 percent in 2036 and 3 in 2046; the middle forty percent is close to break-even, plus 0.1 and plus 0.5; the richest ten percent gains 5.7 percent in 2036 and 13.1 in 2046, and in the best future in twenty exceeds 14 already in 2036. In the emerging economies everyone loses except the richest ten percent, which gains about 0.3: part of the profits is booked elsewhere, and local capital benefits little. In the poor countries the effect starts small, but the poorer half ends up losing 3.8 percent in 2046.

About 84 percent of the world population lives in groups that consume less than in the world without AI. And within each

group a small average loss comes with much larger individual damage.

Demand for goods remains relatively strong in most simulations; in the worst future the shortfall reaches 2.5 percent of output, but in seven futures out of ten it never appears. But holding up the demand for products is one thing, the ability to create jobs another.

And the model also indicates what to watch in the coming years to understand which road we are taking. Towards a harder transition: a fall in hiring that extends beyond people in their thirties, unemployment in exposed territories while GDP grows, a labor share falling by more than half a point a year. Towards inclusion: new tasks visible on a large scale, rising wages in roles that work with AI, tax rules able to follow incomes. Towards slowed adoption: corporate AI spending stalled and flat productivity until 2030. They are checkable: whoever wants to know, five years from now, whether this book was right will know where to look.

Comparing the scenarios now makes it possible to ask which interventions improve results in the largest number of futures, without confusing income protection with the recovery of work. That is the subject of the next chapter.

Chapter 22

Interventions, Responsibilities and Timing

A policy is not judged by its press release. It is judged by the path a person must follow to use it. An available course, support while waiting and a job at the end are three different things; if one is missing, the other two are of little use. This chapter puts nine interventions to the test, then combines them into a package, and finally says who should do what, and when. The timetable comes from the model, which runs to 2046.

Two Questions Before the Ranking

The comparison of policies uses 2,000 pairs of simulations, a different sample from the 4,000 of the previous chapters. Each policy is applied both to the world with AI and to the one without, and the difference says something precise: how much that policy changes the effect of AI. The tables in chapter 13 also keep the other comparison, with the world without AI and without the policy, which contains the redistribution that would happen anyway.

It is a distinction that matters. A tax reform can make the poor better off even without AI, and that benefit must not be attributed to the fight against the effects of the technology. The dividend in chapter 13, for example, still leaves the consumption of the poorer half about 0.4 percent below the world without AI with the same reform: that is the technological disadvantage that remains, separate from general redistribution.

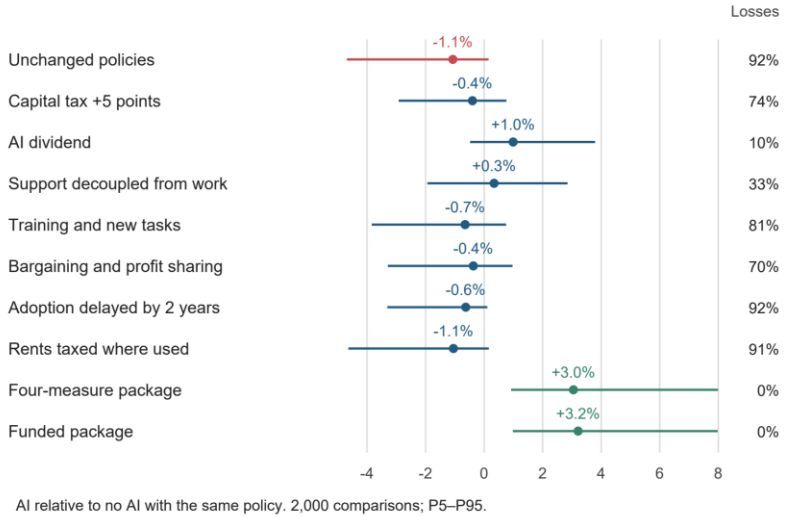


Figure 34 – Nine interventions and unchanged policies. Model simulations: median and P5–P95 of bottom-half consumption in 2036, AI relative to no AI with the same policy. On the right, rounded frequency of losses: 0% for the packages is not a universal guarantee. The reference is not the world without the policy.

The dividend makes the comparison positive for the poorer half in nine futures out of ten; in 2046 the advantage shrinks to about 0.6 percent. And giving income does not create first entries: those who were not hired stay out, even with the dividend. The fund, moreover, works in the model by distributing part of the additional profits; in reality it needs shares, rights and management rules, discussed in chapter 27, and the model does not calculate how much it costs to build.

Training and new tasks reduce additional unemployment. But for consumption of the poorer half the comparison remains negative in 81 percent of cases in 2036, and about 1 percent below in 2046. This lever combines training and new tasks across the whole economy.

Training works when it leads to a specific job. Summaries of programs linked to specific sectors report gains of 14 to 38 percent according to the International Monetary Fund, 11 to 40 according to Katz and co-authors (2022); chapter 31 distinguishes the different versions of the research. They measure the wages of those who took part, not a country's unemployment.

Income support unconnected to work protects the wait. With the same measure, in 2036 the poorer half gains about 0.3 percent, but one future in three keeps a loss; in 2046 the loss returns, about 2 percent. The cost and the financing are those of chapter 13. Supporting those who wait and shortening the wait are two different tasks.

Bargaining and profit sharing raise the pay of those in work and improve the public accounts a little, but in the model they do not create jobs. And since the poorer half lives little on wages, it benefits little. A rule that helps those inside may never reach those who have yet to get in.

Postponing adoption buys time, and pays for it with part of the growth. In the figure the poorer half still loses in 92 futures out of a hundred. The time bought is worth something only if someone knows how to use it.

Finally, attributing AI rents to the country where AI is used shifts revenue towards the emerging and poor economies, and changes the world consumption of the poorest little. But it changes who has the money to help them.

Four Tools Together, and a Promise to Limit

No single policy solves the problem. The package combines them: dividend, income support, training with new tasks, profit sharing. The result is clear. The poorer half moves from a loss of about 1.1 percent to a gain of about 3 percent in 2036 and 3.3 in 2046, com-

pared with the world without AI with the same policies. Losses remain in about 0.2 percent of cases: two futures in a thousand. The rounded zero in the chart should be read as "almost no loss", not as impossibility.

The package's financing uses the assumptions of chapter 13. But the model takes for granted that the state manages to collect taxes, pay out the money and reach people in every territory. In reality each of these three things has to be built.

And the package does not bring unemployment back to the level of the world without AI: the recovery is partial, as in chapter 13. The nine policies mainly test ownership, transfers, training, wages and adoption timing. That is why chapter 24 tests different tools, which act directly on hiring, public employment and working hours. Income and work must be kept on separate lines: winning on one can leave the other incomplete, and requires different tools and checks.

What of all this reaches a life? A service can become easier to use while becoming harder to finance: the two effects must be looked at together, without one cancelling the other. Income support is not a job. A free service does not pay for transport or rent: adjusting consumption for cheaper services softens the average loss, but a quality benefit does not replace a cash expense.

Think of those who live inside these numbers. The twenty-four-year-old graduate, in Europe or North America, who finds fewer open doors in law, accounting, software and marketing. The mid-career employee, left with more supervision and more contact with customers. The nurse, the electrician, those who work in care, supported for now by scarcity. The pensioner in Italy, Japan or Germany, helped by services but with a pension paid from falling contributions. The call center or back-office worker in the Philippines or India, for whom that job was the ladder up. The farmer or clinic worker in sub-Saharan Africa, who can gain from better

information, if there is energy and connectivity. And the owner of capital, who sees returns grow.

They are portraits, to understand where the numbers strike.

Every Action Has Someone Responsible

The table links actions, model levers, those responsible and timing. Finding the money is not enough: you need to know who can act on each step.

Table 43 – Ten interventions, levers, responsibilities and proposed timing.

Intervention	Proposed action	Model lever	Proposed actors and start
Tax capital where AI rents arise	Dividends, capital gains, excess profits; extend the corporate minimum and taxing rights to user countries	Fund transfers and close about 1% of GDP of the ADV gap, 1.2% in C	G20, OECD, national treasuries; 2028
AI dividend funds	Public or sovereign stakes in AI and computing infrastructure, especially with public money; dividend to residents. Precedents: Alaska and Norway	Dividend: strongest single lever on incomes without depending on tax rates	National governments, EU; 2030
Support decoupled from employment	Wage insurance for displaced workers, short-time work with conventional support, floor with negative income tax, benefits portable between employment and self-employment	Support: soften the unemployment peak	National governments; 2030, before the peak

Intervention	Proposed action	Model lever	Proposed actors and start
Protect the first job	Credits or contribution relief for exposed junior hires; mandatory apprenticeships in large firms using AI; public graduate programs	New tasks; continuity of skills	Governments and large employers; 2027, already urgent
Sectoral training funded by adopters	Levies on large AI users for sectoral programs; US evaluations: gains of 14–38% (IMF summary)	Creation of new tasks	Governments, employers, unions; 2028
Invest in sectors with demand and labor	Elder care, health, education, energy transition; ageing sustains demand, AI complements labor	New tasks and absorption of displaced workers	Governments; 2028
Bargain over AI adoption	Consultation, profit sharing linked to productivity, agreements such as the 2023 US screenwriters' contract	Profit sharing	Unions, employers, legislators; 2029
Remove incentives for substitution	Review accelerated depreciation and tax breaks favoring machines over hiring; assess the impact on public services	Slower and more orderly adoption	Treasuries and administrations; 2028
Help for outsourcing countries	Training and digital infrastructure for the Philippines, India and others, funded by a share of global AI revenue	Protect exposed emerging economies	Multilateral banks and G20; 2029
Pensions beyond wage contributions	Part of the funding from general revenue and capital	Protect welfare as the labor share falls	National governments; 2030

One row deserves careful attention: protecting first jobs. It is the most urgent action, and it is not among the nine simulated policies. Credits for hiring young people, compulsory apprenticeships in large firms that use AI, public programs for new graduates require a separate evaluation: who enters, what skills they acquire, what would have happened without, how much it costs. The number for the training lever cannot be copied onto the first-job row.

The Alaska and Norway funds show that a public fund paying dividends can be built; they do not prove that a fund based on AI will yield as much or be distributed in the same way. German short-time work and the 2023 American screenwriters' contract show possible tools, not ready-made instructions for a country's laws.

Governments, firms, unions and international bodies have different powers. Some actions require agreements between countries, others a local decision, others offices able to deliver the help. And an obstacle that is tiny in world totals can decide a person's life: evaluation must choose the right scale, without confusing importance for a person with weight in the totals.

When should the alarm go off? The book proposes three thresholds. Hiring of young people in exposed jobs falls by more than 25 percent, or the fall extends beyond age thirty. The labor share falls by more than half a point a year for two years. Vacancies and hiring in exposed occupations fall for two years while GDP grows. When two thresholds out of three are crossed, support should widen automatically.

Before using them you need to define what is exposed, how hiring is counted and with which sources. An administrative signal arrives before the unemployment rate.

The timetable, if you want to be ready, is this. Between 2026 and 2028, preparation: tax rules, first jobs, training, monitoring. Between 2029 and 2032, construction: portable support, funds, services. Between 2033 and 2036, absorbing the peak: protecting pensions and entries. A system built after the peak arrives late in all futures.

To check that no step has been skipped, the next chapter puts diagrams and equations side by side. Following the flows clarifies where a result depends on an assumed behavior and where on an account that must balance. And a diagram must also declare the arrows it does not draw: its clarity is useful only if it does not hide its limits.

Chapter 23

The Causal Structure of the Model

In 1949 a young New Zealand economist, Bill Phillips, built in a garage in Croydon, south of London, a machine made of tanks, pipes and valves. It was called MONIAC, and it was a hydraulic model of the British economy: colored water represented money, and flowed from incomes to consumption, from taxes to public spending, from savings to investment. Opening a valve, you could literally see where the money went. Several were built, and one is still at the Science Museum in London. Its lesson is simple: in an economy every euro that leaves one place must enter another.

Between "AI automates" and "people are worse off" there are at least five steps: work, incomes, consumption, investment, public resources. Skipping one turns a possibility into a certainty. It is the most common error in the debate on artificial intelligence, in both directions: among catastrophists and optimists alike. This chapter shows the design of the model, to see where each euro goes and where a conclusion depends on an assumption. And it also declares what the design leaves out.

The Circuit That Passes Through the Home

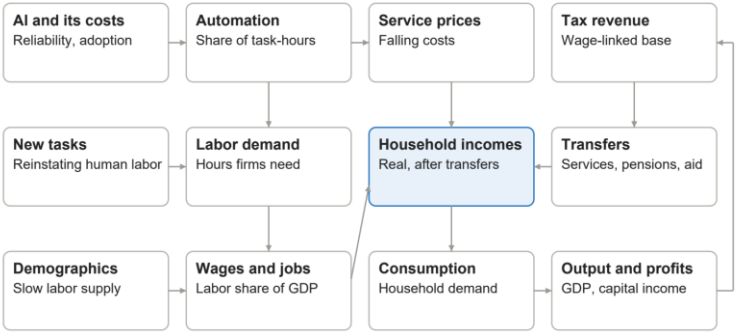


Figure 35 – Diagram of the economic links. Twelve nodes and conditional connections, with no estimated strength or net sign.

Follow the circle. A household's income, after taxes and benefits, decides how much it can spend. Spending becomes demand for goods and services, and demand becomes output and income for firms. Output and profits pay taxes; taxes pay for services and transfers, which return to households. The circuit closes where it began.

A lower wage can weaken this circle. But it does not break it on its own: other incomes count, as do how profits are distributed, how much the state collects, how consumption reacts. The diagram shows the steps through which the spiral can appear, or be stopped. Saving is not consumption, and households with different incomes and wealth do not spend in the same way. Support can keep spending up; an investment by firms can create demand before those looking for work find it.

The "households" in the diagram are not all the same. The model distinguishes the poorer half, the middle forty percent and the richest ten percent, in each of the three regions, and to under-

stand their consumption their different sources of income must be kept. There is also time: automation and new tasks can arrive in different years, and demography changes how many people are available. A region with a stagnant or falling labor supply experiences the transition differently from one full of young people looking for their first job; these differences are part of the calibration.

And there are prices. A lower cost can increase what a household obtains with the same salary. But it depends on how much of the saving reaches the price, on how reachable the service is, on its quality: not all the cost avoided reaches the customer.

Two arrows are missing from the drawing, to keep it readable: from labor income to public revenue, and from AI investment to demand and output, even independently of consumption. They are missing from the drawing, not from the economy the model represents.

Who Pays and Who Receives

The second diagram follows payments, that is, who hands money to whom, as well as how effects are transmitted.

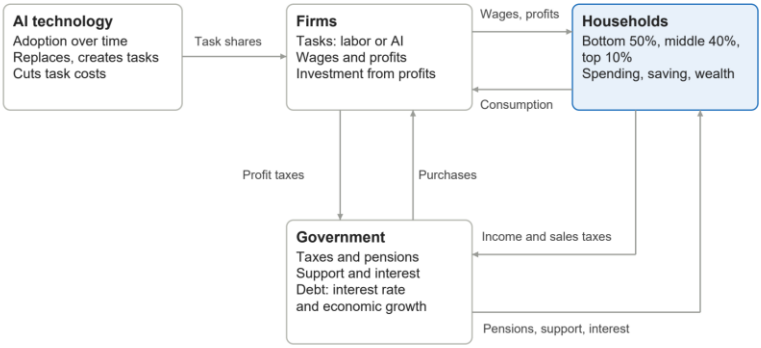


Figure 36 – AI technology, firms, households and the public sector. The diagram shows one region; the model links three through the geography of AI profits.

Firms pay wages and distribute profits to households. Households buy from firms. The state collects taxes on profits, incomes and consumption, then buys goods and services and pays pensions, support and interest. Technology changes tasks and costs, but it does not appear as one more household consuming the wage saved.

Here is the point many forget. A wage saved does not vanish from all the accounts: it becomes a profit, a lower price, an investment, a different product. As in Phillips's machine, the water is not lost. But passing from one pipe to another can change the distribution of income, the share that is spent and the base on which taxes are paid. The money is still there: what has changed is who receives it. And the fact that the accounts balance does not guarantee that every person keeps their income, nor that all productive capacity is used.

The accounting check verifies that every flow has a sender and a recipient, and that the accounts close. Checking behavior is something else: it asks whether the model's rules describe well how

people and firms react. A model can have perfect accounts and get the reactions wrong.

And the use of a technology and the income it produces can be in different regions. The model assigns part of AI profits to the advanced economies; the diagram of a single region does not show the leak, but the leak is there.

The Arrow Has a Memory

Two relationships hold everything up. The first links AI to human tasks: substitution removes them, background erosion reduces them, new tasks add them, but with a delay. The second links AI to growth: savings on automated tasks, help for human work, the contribution of new tasks. Section A10 of the appendix gives both relationships, term by term. To reach jobs, incomes and taxes you then have to continue along the path, through demand, wages and public accounts.

A large saving on a task is not a large saving for the economy. What counts is how many tasks are affected, how much cost is avoided, how much work is created or preserved. And the extra output must find someone to buy it, according to the closure described in chapter 5. From here you get to work: labor content and output decide how many people are needed at the wage considered; supply limits how many can be employed; wages adjust according to the assumptions. But a lower wage does not recreate all the tasks removed. And to get to consumption and public revenue you have to go beyond the growth equation.

The sequence, therefore, does not lead on its own either to poverty or to well-being. It shows where a conclusion depends on an assumption, and what evidence could change it.

Take a café that automates ordering. The time saved can become better table service, fewer paid hours or a higher margin. If

the saving becomes a fee paid to the system's supplier, part of the value leaves the neighborhood. If it becomes wages, it returns in the purchases of nearby households. The café makes the arrows of the diagram visible: the technical saving, the income distributed and the spending that follows are three different steps. The number of orders handled automatically does not say which of the three will happen.

The Questions Left Out

Job quality, freedom of choice, care, learning alongside someone who knows: none of these can be derived from GDP.

The first diagram shows the mechanisms; the second keeps the flows. How people will react to a policy is told by the data, with the check of chapter 18.

The checks in chapters 11 and 12 verify starting points, orders of magnitude and consistency of the simulations. The accounts must close every year, not only at the end.

If a question finds no place in the diagram, there are only two roads: widen the model or narrow the conclusion. Filling the gap with interpretation is not allowed.

One of these gaps concerns learning. The new-tasks parameter explores their timing and breadth, but it does not explain how a person learns to perform them. The next chapter tackles exactly this: the difference between being helped by a machine and knowing how to do something.

Chapter 24

Training and Learning

In Plato's Phaedrus, Socrates tells an Egyptian myth. The god Theuth presents to King Thamus his most precious invention, writing, and calls it a remedy for memory. The king replies that it will do exactly the opposite: those who learn to write will stop exercising their memory, and will have the appearance of wisdom without wisdom. Twenty-four centuries have passed, and writing has not made us stupid. But Thamus's question was not foolish, and it comes back every time a technology does for us something we previously had to learn to do.

With artificial intelligence it comes back in a precise form. Doing a task well with AI and knowing how to do it are two related but different results. A tool can improve today's answer without teaching you to recognize tomorrow's error. To move from training to work, therefore, you need proof of learning, and occasions in which to use the skills learned.

The First Attempt

Almost everyone's first job is made of small tasks: checking a document, preparing a draft, sorting an archive. There you practice with limited responsibility and receive corrections. If those tasks are automated, equivalent occasions for practice must be designed: seeing correct answers does not, on its own, produce experience.

The small task, in reality, always had two products. One was the check done for the firm. The other was the familiarity the beginner acquired with documents, exceptions, people. The first was visible at once in the day's account. The second emerged much later, the day that person recognized a problem they had never seen.

This is not a reason to keep every repetitive task. Practice can be built even when an activity is automated: simulated cases, mentoring, comparing solutions, growing responsibilities. But some-

one has to organize them, devote time to corrections, check that they work. And two measures must be kept separate: the work completed with the tools available, and what a person understands and can transfer to a new case. They can improve together, but it is not guaranteed. An evaluation must know which of the two it is looking at.

In 2023 and 2024 Bastani and co-authors followed almost a thousand mathematics students in a Turkish high school. Classes were assigned to three groups: without AI, with GPT Base, a standard version of GPT, and with GPT Tutor, designed together with the teachers. The results were published in PNAS in 2025.

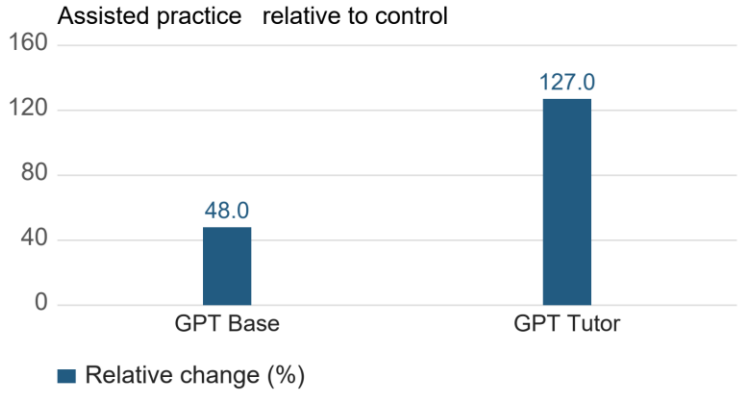


Figure 37 – Bastani and coauthors, PNAS 2025: relative improvement in assisted practice. The chart reports the published summary estimates, without intervals; the exam without assistance is shown separately in Table 44.

During practice sessions with AI, results improve by 48 percent with standard GPT and by 127 percent with the tutor, compared with the group without AI. It looks like a triumph.

Table 44 – Two moments of the same evaluation. Source: Bastani and coauthors, 2025.

Tool	Assisted practice	Exam without assistance
GPT Base	+48% relative	−17% relative
GPT Tutor	+127% relative	Difference not statistically significant

Then comes the exam, without help. Students who had used standard GPT do worse than those who had never used it: minus 17 percent. Those with the tutor show no statistically significant difference from the group without AI. With the tool in hand, the students seemed better. Without it, they were worse. They had done the exercises without learning them: Thamus would have recognized the problem. It is a short-term study, in a specific context; but it is hard to ignore.

To evaluate skills you have to ask how a machine's error is recognized, how a new request is handled, how you go on when the tool is not there. A test without help measures some abilities, a test with help measures others, and a job can require both. And being able to work alone does not mean giving up technology: it means knowing when to trust, when to check, when to ask for help. Receiving a complete solution and explaining a decision before comparing it with an alternative are different teaching activities.

Another school gives another answer. De Simone and coauthors, in World Bank working paper 11125, of 2025, evaluated a six-week program in Nigeria, with an AI tutor and guidance from teachers. The effect is 0.31 standard deviations in the overall assessment and 0.23 in English, the main outcome. These are results of the program in its context, not of the language model alone.

A standard deviation measures the effect relative to how much scores vary. The 0.23 does not mean 23 points out of a hundred,

nor 23 percent more salary. An example translates the unit: if scores varied by ten points, an effect of 0.23 standard deviations would correspond to 2.3 points on that scale.

$$0.23 \times 10 = 2.3\,points\ (example)$$

The two studies do not add up to an "average effect of AI": the minus 17 of one and the 0.23 of the other have different units, subjects and tests. Together they say something else: AI can help people learn, or help them finish the task without learning. In Turkey students received answers; in Nigeria they were guided by teachers. The difference is made by the design, not the tool.

From Know-How to a Job

Even a learned skill must meet someone who will pay for it. Knowing how to do a new activity does not guarantee that an organization has the money or the customers to open a position. In the model the recovery of tasks is represented by ψ and the time by the delay; the age module distinguishes access, skills and the match with demand. Better training helps some of these steps, but it does not create a customer, a firm or a funded position. It must be followed up to hiring, not up to the certificate.

An evaluation track first defines what must be learned, then tests skills on new cases, then follows applications, interviews, entries and duration of work. If the path breaks off, you need to distinguish whether what was missing was learning, transport, an acceptable offer, or simply demand.

Table 45 – Teaching outline for following training and access.

Step	What to observe	What to distinguish
Learning	Test before and after	Assistance available and autonomy

Step	What to observe	What to distinguish
Access	Attendance and dropouts	Outcomes of those who finish and those who leave
Job search	Applications and interviews	Preparation and available opportunities
Entry and duration	Hires and retention	Initial access and stability

The table prevents any recorded attendance from being called a success. If a hundred people start a course and sixty reach the final test, the results of the sixty do not tell the path of the other forty, and often the forty are precisely the ones the course wanted to reach. A free course still costs time, travel, someone to look after the children. And if half the participants find work, you have to ask how many would have found it anyway, and how many enrolled precisely because they were more able to complete the path. That those who attend work more does not prove the course works.

Three Ways to Open the Door

Chapters 13 and 22 showed that income can be protected. But work? The book ran an experiment that acts directly on jobs, in the age module, with three tools: public employment, an incentive to hire young people, reduced working hours. Same budget for all: 0.7 percent of each region's output every year, from 2027. Four thousand combinations of assumptions.

The first tool creates public jobs: the budget divided by the annual wage gives the number of additional jobs, distributed across ages according to potential supply. These are available jobs that still have to meet people. And the skills required, the organization

of services and material resources are left out: paying a salary does not by itself build a clinic.

The second is an incentive to hire people aged 15 to 34, equal to half the annual wage. With a trap known to anyone who has studied these policies: some of the hiring would have happened anyway. The model assumes it is between 40 and 80 percent. If six hires out of ten were already decided, the incentive pays for ten people but adds only four. Counting beneficiaries does not measure the effect.

The third reduces working hours. The same labor requirement is divided among more people, without new output. The state makes up 60 percent of the wages lost for the hours cut; the worker loses the remaining 40, and the hourly wage stays on the macroeconomic path. With the same total hours, shorter hours can put more people to work; how many depends on the budget, the share of wages made up and the labor share of output. It works if the work can be divided: shifts, skills and coordination can prevent it. Section A11 of the appendix shows the calculation. In 2036 the average cut in hours in the advanced economies is about 2.35 percent, and those already employed lose about 0.94 percent of their labor income. The aggregate figure does not say whether a given household can afford it.

Table 46 – Effects of the three policies in the world with AI, relative to the same world without intervention. ADV, reference scenario, 2036, medians in percentage points. Employment: employed / potential supply of the age group. First entry: hired / eligible F pool, ages 15–24. Total non-employment: 1 − employed / potential supply, including the discouraged.

Intervention	Entry 15–24	Empl. 25–34	Total non-empl.
Public jobs	+1.28	+1.09	−1.16
Youth incentive	+3.54	+2.93	−0.97

Intervention	Entry 15–24	Empl. 25–34	Total non-empl.
Shorter hours	+1.93	+1.79	−1.82

To get young people in, the incentive wins: plus 3.54 points in first entry, against 1.93 for reduced hours and 1.28 for public jobs. For employment of 25–34-year-olds too, the incentive is the strongest. But the last column changes the ranking. If you care about how many people, of all ages and including the discouraged, remain without work, reduced hours win: minus 1.82 points, against 1.16 for public jobs and 0.97 for the incentive. The measure you choose chooses the policy. There is no "best" in the abstract.

The figure shows the ranges for first entry both in the reference scenario and in the one in which retraining is slow: the ranking too depends on the context.

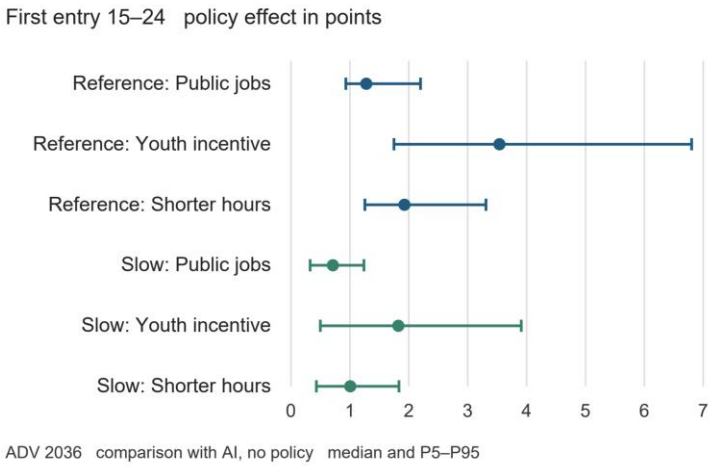


Figure 38 – First entry 15–24 in the advanced economies, 2036. Effect of the policy in the world with AI relative to unchanged policies. Dot: median; bar: P5–P95 of 4,000 draws. Blue: reference; green: slow retraining. Annual budget assumed equal to 0.7% of regional GDP.

And now the figure that weighs most of all. In the code the world without AI also receives each policy, so you can ask how much technological gap remains after the intervention. With the incentive, first entry of young people remains about 11.88 points below the world without AI with the same incentive. The policy helps, and the hole remains: both things are true at once.

The budget is the same for the three policies, and the comparison is between them: the costs have not yet been fed back into the taxes and debt of the macroeconomic model.

There remains the time of those who teach. Mentoring a beginner costs: an expert loses time, and for a while the tasks completed fall. But that cost builds tomorrow's capacity. A firm can look only at the immediate cost and the tasks saved, or also at the future capacity of those who learn. Whether it pays depends on how much that capacity will be needed, on the ability to retain trained people, on who takes the benefits. If training produces advantages outside the firm too, those who pay today may not recover them: the trained young person can go to the competitor. It is an economic reason to discuss who finances learning.

And if many firms eliminate entry-level activities and look only for experienced staff, who will train tomorrow's experts? Each firm, alone, makes a convenient choice; all together they drain the reservoir. The risk exists even when every single decision seems reasonable. To prepare, different entry routes can be designed and their results checked: explaining a decision before seeing the machine's answer, practicing on errors in a protected environment, having a contact person for doubts.

Imagine a library in which a system prepares the first record for new volumes. A newcomer can check authors, dates and classifications, compare the proposal with the sources, explain the corrections. Or they can receive a single order: approve, quickly. In the second case the record is completed, but the learning is not visible.

A test on a different volume, first without the tool and then with it, says which abilities have been acquired.

Readable training links what it teaches, the test that checks it and the work in which it will be needed. But getting in is not enough. After entry, hourly pay, hours, purchasing power and stability count. The next chapter separates these dimensions, and shows why a job found again can be a worse job.

Chapter 25

Job Quality and Pay

On 31 December 2020 the Court of Bologna filed a ruling that went around Europe. It concerned the algorithm with which a large home-delivery platform assigned shifts to its riders. The system rewarded those who showed up for booked shifts and penalized those who did not, without asking why: whether the rider was ill, had a child to look after or was taking part in a strike, it was all the same to the algorithm. The judge found it discriminatory, because it equally penalized absences due to reasons protected by law, such as a strike or illness: it treated different situations as the same.

The statistics will say you have a job. They will not say what time the shift ends, who decides it, whether you can contest it, whether you earn enough to live on. An offer with a higher hourly wage can give less income per month, if it offers fewer hours. And even returning to work can leave conditions worse than before. To read a re-employment you have to separate entry, income, continuity and the ability to organize your own life.

Two Offers, a Misleading Number

Take a person who earns twelve euros an hour for a hundred and sixty hours a month. They lose their job and receive two offers. The first: thirteen euros an hour, but only a hundred and twenty hours. The second: eleven euros eighty, for a hundred and sixty hours. Looking only at the hourly wage would hide half the result.

$$Monthly\ income\ =\ hourly\ wage\ \times\ hours\ paid$$

Table 47 – Example of hourly wage and gross monthly income. No taxes, transfers, transport costs or price changes are included in the calculation.

Case	Euros/hour	Hours paid	Euros/month
Before	12.00	160	1,920
Return A	13.00	120	1,560
Return B	11.80	160	1,888

Before, they earned 1,920 euros gross a month. The first offer pays more for each hour and brings 1,560 euros: 360 less. The second pays less per hour and brings 1,888 euros, almost as before. Those who look only at the hourly wage choose the first and get it wrong. But to compare them with a person's life you need other information: what they prefer, what constraints they have, what contributions they accrue, how predictable the hours are, how much it costs to get there. The formula clarifies a measure; it does not choose for anyone.



Figure 39 – Teaching indices calculated from Table 47, initial situation equal to one hundred. Hourly wage and monthly income can move differently. Example, not a model outcome or an estimate of an observed loss.

And reduced hours can be wanted, or imposed by the lack of alternatives. The same number of hours does not say which of the

two cases you are looking at. Even someone with a job may not have the hours they would like, and the statistic that counts the employed does not distinguish those who work little because they want to from those who work little because they cannot find more. Wages, hours and freedom of choice require separate measures.

Then there is purchasing power, which is not written in the job advertisement. A wage is measured in what it buys, and the difference between income in euros and purchasing power remains necessary even when pay rises. If pay rises from 100 to 110 and prices from 100 to 115:

$$Real\ index = (110 / 115) \times 100 \approx 95.65$$

Real pay falls to about 95.65. A raise on paper, a loss in the shopping cart. And time counts as much as prices: two jobs with the same net salary can require half an hour or two hours of travel a day. And compatibility with transport can decide whether a job is reachable. That is why the book keeps the dimensions separate, instead of adding them up into a single index of the "good job".

A Job Has More Than One Quality

The OECD, in a 2015 paper by Cazes, Hijzen and Saint-Martin, divides job quality into three parts. Earnings quality: how much is earned and how pay is distributed. Labor market security: the risk of losing the job and the economic consequences if it happens. Quality of the working environment: the conditions in which people work. The book uses these dimensions to order the questions.

Table 48 – Questions for reading the conditions of the new job. The three dimensions follow the OECD framework; examples and checks are proposals for observation.

Dimension	Question in the story	Proposed check
Earnings quality	How much income arrives?	Pay, hours and continuity
Labor market security	How fragile is the return?	Duration, interruptions and support
Working environment	What is the shift like?	Workloads, autonomy and mentoring

Continuity requires knowing how hours are decided and communicated. A sum received once and a predictable salary are not the same thing, even if in one month they coincide. Hours communicated a day in advance are not hours on which to build a life.

The model measures how many people work, not what their work is like. And the unemployment that remains in 2046 says how many have not been reabsorbed, not how those who returned are doing: even if it went back to zero, the second question would remain open.

Not everything, however, goes in one direction. The OECD surveys published by Lane, Williams and Broecke in 2023 collected the views of 5,334 workers and 2,053 firms in manufacturing and finance, in seven countries, interviewed at the start of 2022. About 80 percent of those who used AI said their performance had improved; 8 percent said the opposite. These are perceptions collected before the new generation of chatbots, but they must be stated. A system can remove dangerous or repetitive activities and improve work. It can also speed up its pace. Freed time does not decide on its own where it goes: to production, to breaks, to pay, or to a heavier workload. You have to look at how the benefit is divided.

When the Program Assigns the Shift

There is a way of using data that already concerns millions of people: algorithmic management. A program assigns tasks, predicts times, evaluates performance. To understand its effects you have to look both at the usefulness of the information and at the rule that turns it into decisions.

The ILO distinguishes algorithmic management from generative AI: organizing and evaluating work with data-based systems can include rules that learn nothing, and merely apply. A joint study by the European Joint Research Centre and the ILO, by Rani, Pesole and González Vázquez, published in 2024, examined cases in logistics and health care in Italy, France, India and South Africa. It found efficiency gains, and at the same time problems for working conditions and surveillance.

The same instruction can be a help, or a constraint that ignores important circumstances. Describing the system is not enough: you need to know how it is used and what can be corrected. Think of a system that flags a delay compared with the expected time. The manager can ask what happened. Or they can treat the flag as definitive proof of poor performance. The measure is the same; what changes is the possibility of explaining and correcting. That is exactly what was missing in Bologna. To be accountable, an organization must make explicit the goals, the events the system does not see and who has the power to review a decision. A score does not describe a job, and it does not replace checking an error.

The Number of Those Who Came Back

Ten people returning to work are not necessarily the ten who lost it, and the number does not say in what conditions they returned. To evaluate people you need data that follow them over time. A record of this kind registers the previous job, the time without

work, the next job and comparable measures of quality; it distinguishes direct moves, returns and first entries, and also follows those who drop out. And to employment, hires, separations and wages it adds hours and conditions, when the data allow.

Looking only at those who accepted a new job would leave out those who did not find one. And there is a frequent statistical trap: the average pay of the re-employed can rise even if nobody is better off, because those who returned were mainly the luckiest. Those who stay out disappear from the average. Those who lose a job and those who find one are not necessarily the same people: you have to keep the time between one state and the next, and the information on who stays out.

Turning the time saved by AI into shorter hours at full pay is an idea that comes up often. It works only if the gains are really there, if they are distributed and if the business holds up. Reducing only the paid hours, as in the opening example, reduces income. Reduced hours, hiring incentives and public employment must be studied: the experiment in chapter 24 compares some versions with the same budget. A serious experiment should ask whether the time saved really becomes free time, whether people keep their income, whether the service keeps its quality. And it should look at intensity: doing the same work in fewer hours can mean working faster.

An offer can be good in the pay account and impossible in the account of the day. If the shift ends after the last bus, getting home requires another solution and another expense. Hours communicated at short notice can make even an already booked nursery place unusable. The pay table remains correct, but it tells only part of the comparison. The case says what information to add: actual hours, connections, costs, predictability.

Actual hours, shifts, pay, continuity, mentoring: to understand why they change you have to look inside the firm, at its costs, its

customers, its organization, its suppliers. That is the subject of the next chapter.

Chapter 26

AI Adoption in Small Firms

In the big studies, artificial intelligence enters multinationals. But in Italy about 95 percent of firms have fewer than ten employees, according to Istat: workshops, professional practices, shops, small garages, laboratories. There AI does not arrive with a strategic plan and a team of consultants. It arrives with a thirty-euro monthly subscription. And the bill is different from what it seems.

Generating a draft, checking it, producing the good, delivering it: these are different phases. A small firm can save time in one phase without reducing staff, or keep everyone and buy fewer outside services. To measure the effect of AI you have to follow the work inside the firm and outside it.

Firms Do Not Start Together

If AI speeds up one phase, the bottleneck moves to the next. The overall benefit depends on checking, preparation and demand: the technical gain does not automatically become output sold.

And firms do not start together. According to Eurostat, in 2025 at least one AI technology was used by 17 percent of small European firms, 30.36 percent of medium-sized ones and 55.03 percent of large ones. The survey concerns firms with at least ten employees in the sectors indicated, and it measures adoption: it does not prove that growing makes firms use AI.

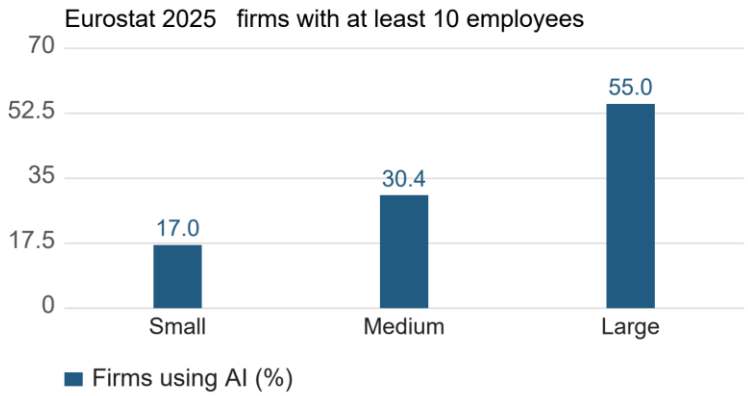


Figure 40 – Eurostat, 2025: firms using at least one AI technology, by size. Small: 10–49 employees; medium: 50–249; large: at least 250. Percentages of firms, not of jobs replaced. The figure excludes firms with fewer than ten employees.

The large ones run, the small ones follow. And microenterprises, those under ten employees, that is, most Italian firms, are not in the survey. It does not mean they do not use AI: it means the sample does not see them.

The subscription, moreover, is only the first item on the bill. There is the time to learn to use the tool, there are the attempts to throw away, there is the checking of what it produces. Do the sums: a fee of thirty euros a month, two hours of checking valued at twenty-five euros each, five euros of other expenses. That makes eighty-five euros every month. Plus twelve initial hours of preparation, three hundred euros, once only.

Table 49 – Teaching account of an adoption. The value of time is an opportunity-cost assumption: it is not an observed wage or money automatically collected.

Item	Amount	When
Fee	30 euros	Every month
Checking: 2 hours × 25	50 euros	Every month
Other expenses	5 euros	Every month
Recurring total	85 euros	Every month
Preparation: 12 hours × 25	300 euros	Once
Benefit: 4 hours × 25	100 euros	Every month, by assumption

If the program frees four hours a month, valued in the same way, the benefit is a hundred euros. Minus the eighty-five of recurring costs, the balance is fifteen. But the first month, counting preparation, is minus two hundred and eighty-five. A low price can hide a high initial cost.

$$Monthly\ balance\ =\ h\,v\ -\ C$$

Here h is the hours freed in the month, v their assumed value, C the recurring cost. With h equal to four, v equal to twenty-five and C equal to eighty-five, the balance is fifteen euros; the initial cost stays out of the monthly formula and must be counted separately. Now change a single assumption: checking requires four hours instead of two. Checking costs a hundred, plus thirty for the fee and five for other expenses: a hundred and thirty-five. The balance becomes minus thirty-five. Changing one assumption is enough to change the sign of the result.

And that freed time, what does it become? It can become one more delivery, a break, shorter hours, a wait. Calling it revenue at

once would hide the step between capacity and sales. Calling it useless would hide the value of a less hectic day.

Because a proposal is not an order. If AI lets you prepare ten variants instead of three, but the customers stay the same, the time saved is consumed in variants, and deliveries do not increase. If the variants bring new customers, output grows. The difference between productivity and demand also holds for a workshop: reducing the hours for a task changes employment depending on how much the quantity demanded increases. The time saved and the customer to be won are two parts of the same account. And the firm can pass the saving on to customers, keep it as margin or invest it, while a competitor may lower prices first. One firm can earn more without hiring, another hire while earning less.

What Firms Say

Another source looks at the same firms from another side. An OECD study published in 2025 surveyed, in 2024, 5,232 firms with up to 249 employees in seven countries, including sole proprietorships; Italy is not among them. Among those using generative AI, 83 percent said their staffing needs had remained the same, 9.1 percent that they had decreased, 5.5 percent that they had increased. These are firms' answers, not jobs lost or created.

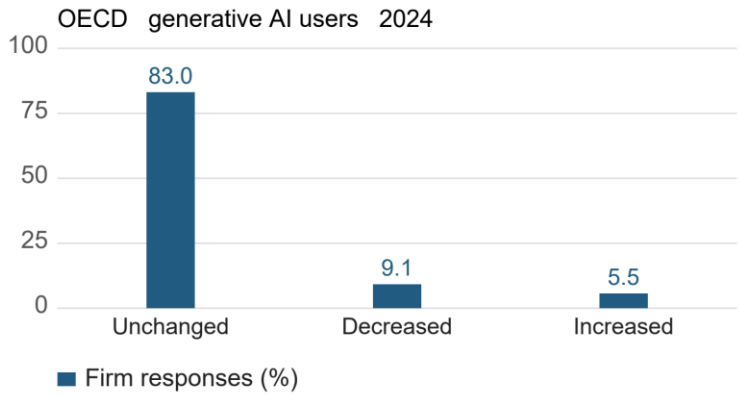


Figure 41 – OECD, 2024 survey: staffing needs reported by firms using generative AI. The three responses shown do not exhaust the total.

Be careful how you read it. "Unchanged needs" does not mean unchanged tasks: the work may have changed even if the people are the same. And the three columns do not measure how much: a firm that says "increased" may have hired one person or ten, one that says "decreased" may have let go one or many. Without the right quantities and weights, subtracting the percentages does not give a net balance of jobs.

Table 50 – Two surveys answer different questions. The comparison clarifies the scope of the sources; it does not build a common time series.

Source and year of data	Scope	Subject
Eurostat 2025	EU; at least 10 employees	Use of AI technologies
OECD 2024	Seven countries; up to 249 employees, including sole proprietors	Generative AI and reported staffing needs

The two surveys answer different questions, and they do not form a common time series. And they describe the first years of adoption.

Then there is an effect that company statistics do not see. A firm that used to pay an outside graphic designer for drafts now does them in-house with AI. Its employees stay the same. But the designer has lost a client. Internal staff does not change, and an income disappears all the same, outside the door. The boundary of the firm changes the reading of the data: a fall in internal staff can come with more work contracted out, and vice versa. Following suppliers too makes it possible to distinguish the two movements.

Employees, collaborators, the self-employed and suppliers are in different accounts. Moving an activity between these forms changes continuity, protections and risks, even if total work remains similar. A career can include intermittent assignments as well as a job as an employee, and self-employment can be freedom, or the way the risk of orders is offloaded onto a single person: the label does not say which. And entrusting a phase to an outside party, or to the machine, does not eliminate responsibility for the result, which for the customer remains with the firm that delivers. Outsourcing, automating or doing in-house can redistribute time without removing the need to check. That is why tasks and relationships must be followed, as well as job titles.

In the model, moreover, AI adoption follows a timetable decided from outside. It is useful for comparing scenarios, but it does not explain why one firm adopts in January and another postpones for two years. In real firms deadlines, liquidity, expected orders and delivery capacity count. And increasing productive capacity does not by itself create customers.

One small firm that saves does not change the economy. A hundred thousand doing it together do. If all of them cut the same expense, whoever sold that service loses orders. If all of them increase capacity in the same market, not all will find the extra customers; if they open different markets, the expansion can be wider.

The Customer Who Postpones

And then there is the circle. If households lose income, they buy less, and pass the difficulty on to other firms. If firms sell less, they hire less.

The model already links incomes and purchases. But in the first version the age module calculated who remained without work and did not return that number to the macroeconomic accounts. The new experiment closes the circle. It returns a measure of employment to wages and benefits; the model recalculates demand, the module recalculates how many people work, and the two parts update each other until the path changes by less than a set tolerance. The world without AI follows the same procedure too, so that the comparison has the same rules.

The new account uses the share of people without work over potential supply, the discouraged included. That share weighs on wages and on benefit spending: those who stop looking continue to count in both channels.

The experiment uses four thousand draws with the same initial assumptions and changes only the link between the parts. Each column compares the world with AI and the one without, with the same architecture. The results concern regions and groups, not a workshop's orders.

Table 51 – Experiment with the two-way link. World bottom-half consumption in 2036 in the three scenarios; the other measures concern the reference scenario. Medians of 4,000 draws. Each column compares AI and no AI with the same architecture; units and territory are indicated in the rows.

Measure	**One-way**	Two-way
World bottom-half consumption 2036 (%) – reference	−1.03	−1.39

Measure	One-way	Two-way
World bottom-half consumption 2036 (%) – slow retraining	−1.43	−3.15
World bottom-half consumption 2036 (%) – persistent substitution	−2.64	−7.24
Wages, advanced economies 2046 (%)	−9.37	−9.87
Primary balance, advanced economies 2036 (points of GDP)	−0.99	−0.97
Debt/GDP, advanced economies 2046 (points)	+11.19	+11.39

In the reference scenario the change looks small: the consumption loss of the world's poorer half in 2036 moves from 1.03 to 1.39 percent.

But look at the difficult scenarios. If retraining is slow, the loss moves from 1.43 to 3.15 percent. If substitution persists, from 2.64 to 7.24. In both cases it more than doubles. The small change in the reference scenario does not allow the conclusion that the circle matters little: precisely when getting back in is hard, or substitution persists, feeding unemployment back into wages and benefits multiplies the damage. Every job lost removes income, the income lost removes demand, the demand lost removes other jobs.

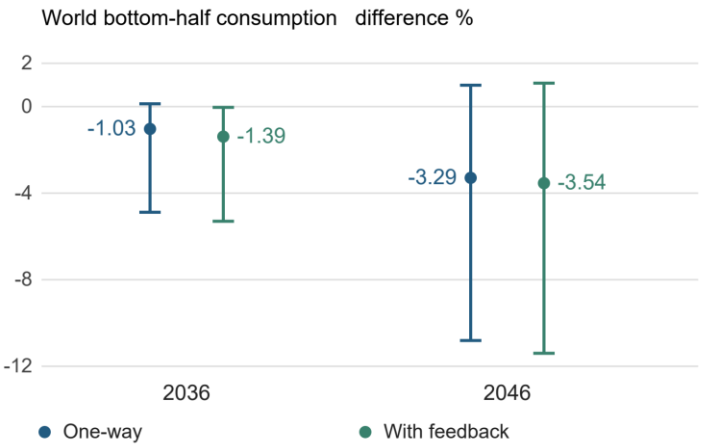


Figure 42 – World bottom-half consumption, AI minus no AI, in 2036 and 2046. Median points and 5th–95th percentile ranges of the reference scenario, before and after the feedback. The two years are shown separately. Source: experiment on the link, 4,000 paired draws.

Not everything gets worse, and it must be said. The debt of the advanced economies in 2046 changes little: from 11.19 to 11.39 additional points. The 2036 primary balance improves slightly: more people without work increase the benefits to be paid, but lower wages reduce the amount of each. Closing the circle does not make every number worse.

In 2046 the consumption of the world's poorer half remains below the world without AI: the median moves from minus 3.29 to minus 3.54 percent. In all scenarios the calculation reached the tolerance, and the accounting discrepancy remained below a thousandth of a billionth of output. The results are numerically consistent.

Take a workshop that makes signs and uses AI to prepare drafts faster. The work goes on: checking measurements, choosing materials, printing, mounting. If orders stay the same, the hours freed in design do not become turnover: they reduce a cost, or remain empty. If new orders arrive, the limit becomes printing or mount-

ing. The cost account in this chapter helps to distinguish the gross saving from the real advantage, after fee, checking and learning.

And the workshop is not alone. If all the workshops in the city adopt the same tool, drafts become faster everywhere, and each one's advantage over the others disappears. At that point the saving passes to customers in the form of lower prices, or turns into fees paid to the software supplier, which now collects from all of them. The graphic designer the workshop once paid for drafts, meanwhile, has lost not one client but ten. Every decision, taken one by one, was sensible. The overall effect is something else, and that is why the book follows both the single firm and the whole economy.

Firms of different sizes take part in the same economy, with different timing. The next chapter follows the money to those who receive it, and distinguishes who owns, who collects and who decides.

Chapter 27

Capital Ownership and Concentration

In 2013 a French economist, Thomas Piketty, published a book hundreds of pages long on the history of wealth, Capital in the Twenty-First Century, full of tables and historical series. It became a worldwide bestseller, something that almost never happens to an economics book. Its thesis, summed up in an inequality between two letters, was that when the return on capital exceeds the growth of the economy for a long time, wealth tends to concentrate in a few hands, because those who own accumulate faster than everyone else's incomes grow. The thesis has been debated and contested. But it put back at the center a question that artificial intelligence makes urgent: who owns the machines?

When a machine does a person's work, the value produced does not disappear. It goes to someone. The decisive question is not how much will be produced, but to whom it will go. An activity can combine land, equipment, software and services belonging to different parties; to understand who benefits from productivity you have to reconstruct the rights to income, not just the place where production happens. The examples in this chapter hold output fixed and change only distribution, to make this step visible.

The Machine and the Right to Income

Those who work receive a wage; those who own receive profits, dividends, rents. Many receive a little of both. But the proportions matter. A small holding in a fund can supplement a salary, not replace it; a large share can support a household even without work.

In the model the population is divided into three groups: the poorer half, the middle forty percent, the richest ten percent. Chapter 19 recalled that the world's richest ten percent own about 75 percent of wealth, the poorer half about 2. When income

moves from labor to capital, it moves towards those who already own.

And wealth is not income. Wealth is what is owned at a moment, net of debts when we speak of net wealth; income is a flow over a period. Not every asset pays a dividend every year, and not every price increase immediately becomes money to spend. A house or a share can be worth a lot and not pay the month's shopping; a firm can make profits without distributing them. Ownership, income accrued and money available are different measures.

Take a hundred people, divided into the book's three groups. A hundred units of capital income, distributed in full. Output, wages and employment stay fixed; only the division changes.

Table 52 – Two distributions of the same capital income. The group sizes follow the book; shares and amounts are for teaching, not data from the original calibration.

Group	People	Income A	Income B
Bottom half	50	20	5
Middle	40	30	15
Top	10	50	80
Total	100	100	100

In case A the poorer half receives 20 units, the middle 30, the top 50. In case B, 5, 15 and 80. The total is the same, a hundred. The world is different.

$$Dg = pg\,\Pi$$

The formula says one thing only: in the example, in which everything is distributed, the income D of group g is its share p multiplied by the total Π, and the shares add up to one. And it separates

the two things that matter. The total is decided by technology. The share is decided by property rights.

Divided per person, in case A each member of the poorer half receives on average 0.4 units, each in the middle 0.75, each at the top 5. In case B, 0.1, 0.375 and 8. Eighty times as much between the first and the last.

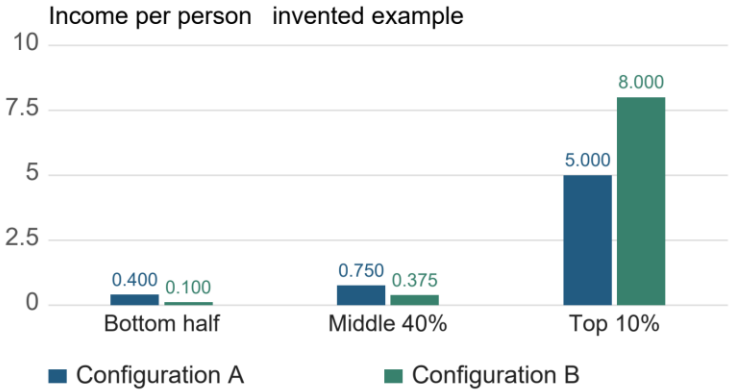


Figure 43 – Average income per person in the two distributions of Table 52. Same total distributed.

The chart makes visible the difference between a group's share and the average amount per person. And it says something worth keeping in mind. The productivity of a technology does not decide on its own who gets its benefits: rights to value come into play afterwards, when capacity has changed and the weight of labor has shrunk. Confusing the two levels makes a distribution that was in fact assumed look inevitable. Those who say that "AI will make everyone richer" are making an assumption about ownership, without declaring it.

Where It Is Produced and Where It Is Collected

In the model the parameter ζ represents the part of AI income that goes to owners or suppliers in other regions. Even before AI

a territory could host production and send some income elsewhere; AI lengthens these routes.

The diagram distinguishes the place of production, the remuneration of labor, the income from goods and services supplied, and the residence of those who receive it.

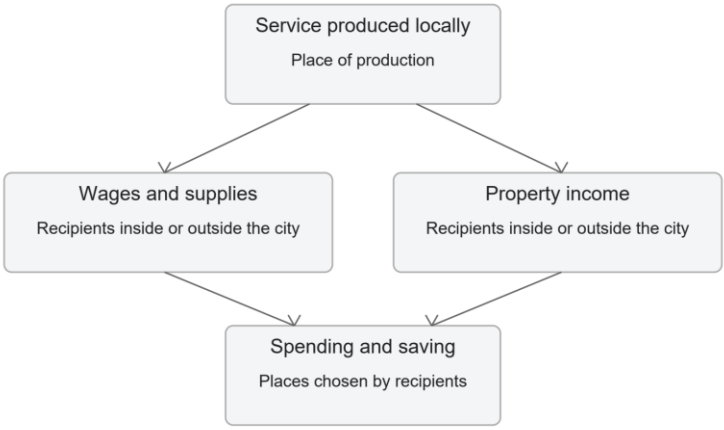


Figure 44 – Diagram for reading the recipients of value and the places of spending.

A service bought from a distant supplier improves the work of those who use it, and moves a payment far away. If you look only at the invoice, you lose the benefit; if you look only at the time saved, you lose who collects. And the territory does not coincide with the registered office written on a document: production, ownership, place of registration and residence of those who receive the income are different pieces of information.

Then there are the firms inside the average. Autor, Dorn, Katz, Patterson and Van Reenen showed in 2020, in the Quarterly Journal of Economics, after a first working paper in 2017, that the American labor share fell partly because dominant firms grew, the "superstars", which pay labor a smaller part of the value they produce. Activity moved towards firms with less labor per unit of val-

ue, and this shift explains part of the overall decline. It is a story of market structure, before generative AI.

An example shows the mechanism. Two firms: one gives labor 60 out of 100 of value added, the other 30. If they weigh the same, the average is 45. If the second comes to produce three quarters of the total, the average falls to 37.5, without either of them having changed its own rules. And AI has exactly the shape of a technology that makes those who own it grow.

To understand whether a market is concentrated, counting the sellers is not enough. Four firms with 25 percent each, and four firms with 70, 10, 10 and 10: same number of sellers, very different markets. To measure it an index called HHI is used: the sum of the squares of the shares. With percentages from zero to a hundred the maximum is ten thousand; with four equal shares it is 2,500, in the second market 5,200. Using shares between zero and one would change the scale, not the comparison. The calculation is in section A12 of the appendix.

Table 53 – Two markets with four sellers. The shares refer to the same hypothetical market and the same period.

Market	Percentage shares	HHI
A	25; 25; 25; 25	2,500
B	70; 10; 10; 10	5,200

A higher index describes more concentrated sales. A market can become concentrated because a firm offers a better or cheaper product; or because it is hard to switch supplier or to compete. The index does not say which of the two explanations is right. Switching AI supplier, moreover, can cost time, training and compatibility, as well as the price of the new service. These are costs that tie the customer to the supplier.

Share, Amount, Price

To obtain the overall labor share, each firm weighs by the value added it produces: labor incomes are summed and divided by total value added. Section A12 gives the relationship. Here is a paradox worth keeping in mind. Two firms, a hundred units of value each, and labor receives 90 out of 200. If the firm that pays labor less, the one with a share of 30, grows to three hundred and the other stays at a hundred, labor receives 150 out of 400. The share falls to 37.5 percent. But the amount paid to workers rises from 90 to 150.

Share and amount answer different questions. And for households what also counts is how many people, which prices, how many hours, how income is distributed within the groups. Output growth does not decide on its own how it is divided.

A more efficient firm can keep the saving, or lower prices and sell more. With a cost of eight and a price of ten, the margin is two per unit, before other costs. If the cost falls to six and the price to nine, the margin rises to three and the customer pays less. If the price falls to seven, the margin falls to one. The fact that the customer gains does not decide, on its own, how value is divided between labor and ownership. How much of the saving reaches prices depends on the alternatives and the possibility of switching supplier.

Taking Part and Deciding

Receiving a dividend, voting at a shareholders' meeting, deciding shifts: these are three different powers. A minimal holding can give an income without giving a voice on investment or organization. Who finances, who receives the income and who influences decisions are distinct questions. More widespread ownership can distribute proceeds without guaranteeing work to those who take part: a dividend has value, but it is not a job.

Ownership, transfers and services are different routes through which value reaches households, to be compared by cost, risk, access and timing. Broadening ownership, with public funds or worker participation, is the route the model rewards most for the poorer half. But it has to be built: where the shares come from, how much it costs to acquire them, who manages them, who bears the losses. The model distributes profits with one line; reality requires institutions.

Take a library that buys an automated search and translation service. Users are better off. But the contract pays a distant supplier. The benefit of the service and the income that leaves are both real, and they follow different routes. The café next to the library may have more customers if visits grow, or notice nothing. Knowing where a technology is used does not say where the value ends up: ownership, contracts and subsequent spending link productivity to distribution.

And having the money is not enough to find a service. Payment makes a service possible only if the capacity, the people and the organization to provide it exist. The next chapter looks at this limit in the sector where it weighs most: health.

Chapter 28

AI and Health Services

If there is one sector in which artificial intelligence promises to do good, it is health. Reading thousands of images without tiring, finding the tumor the human eye does not see, freeing doctors and nurses from bureaucracy. Some of these promises have already been put to the test. The results are real, and they are more precise, and more limited, than the announcements say.

Booking a service, receiving it and benefiting from it are different phases, and AI can improve one without untying the knots of the others. This book's macroeconomic model does not measure health: to talk about it you need specific evidence, and separate attention to activity, quality and access.

Services and People

A care pathway includes booking, travel, meeting the professional, assessment, treatment and follow-up. Counting a single stage can make a service seem complete when it is not. And quantity, quality and access must be kept separate. More tests do not by themselves prove more health; a better diagnosis is of no use to someone who lives two hours from the hospital and has no car. Reducing fatigue and recognizing a problem is worth a great deal, and the technology does it in precise activities.

In Sweden the MASAI study compared two ways of reading screening mammograms: traditional double reading, by two radiologists, and a protocol in which AI routes images to one or two readings and assists the doctors. The women involved numbered 80,033, randomly assigned. The results of the first safety analysis were published by Lång and co-authors in 2023.

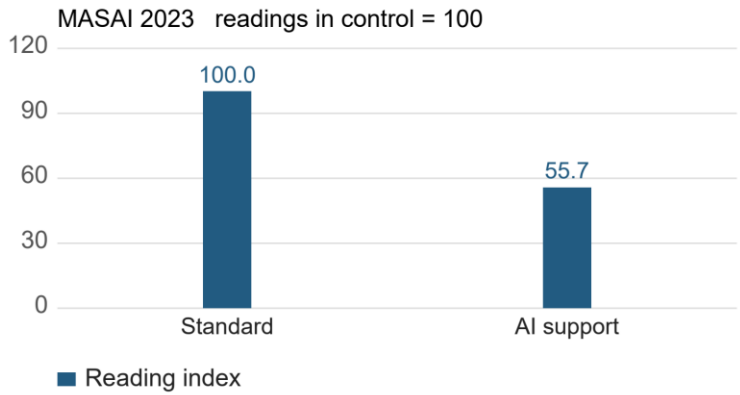


Figure 45 – MASAI, 2023 analysis: index of the total number of readings, control equal to one hundred. The value 55.7 derives from the published reduction of 44.3%.

The readings needed fell by 44.3 percent: taking the reading workload of the traditional method as a hundred, with AI 55.7 was needed. Almost half the reading work, in that protocol, was no longer needed.

And what does that time become? If there are patients waiting, it can serve more people. It can go to other tasks. Or it can accompany a reduction in staff. The chart does not say which road was chosen. But it is exactly the kind of number that, in a budget under pressure, becomes a staff cut.

The study did not stop there. In 2026 Gommers and co-authors published in The Lancet the figure that really counts: interval cancers, those diagnosed between one screening and the next, that is, the ones that were missed. The ratio between the two groups is 0.88, with a 95 percent confidence interval between 0.65 and 1.18. The AI protocol meets the non-inferiority criterion set in advance: it is not worse than the traditional one beyond the established margin. But that interval does not, on its own, prove a certain reduction in missed cancers. Sensitivity, on the other hand, improves significantly: from 73.8 to 80.5 percent, with the same specificity, 98.5, in both groups.

A non-inferiority test checks that the new method does not worsen the result beyond a margin decided in advance; proving an improvement is a different question, and a result can answer the first without answering the second. And some results take years. Less work, more cancers found at screening, no worsening in missed cancers: a serious result, which took years of waiting to be verified.

Another experiment curbs the enthusiasm. Goh and co-authors, in JAMA Network Open in 2024, randomly assigned fifty doctors to two groups: one with access to GPT-4 in addition to conventional resources, the other with conventional resources only, on simulated clinical cases. Median diagnostic reasoning score: 76 percent against 74. The adjusted difference is two points, with an interval from minus 4 to plus 8. No statistically significant improvement. And it is not a measure of outcomes on real patients.

Table 54 – Three evaluations with different questions. Sources: MASAI 2023 and 2026; Goh and coauthors 2024.

Study	Outcome measured	Result and boundary
MASAI 2023	Number of readings	−44.3%; not employment or mortality
MASAI 2026	Interval cancers	Ratio 0.88; 95% CI 0.65–1.18; non-inferiority, margin 20%
Goh 2024	Reasoning on simulated cases	Adjusted difference 2 points; 95% CI −4–8; not significant

Two trials, two different results, because they concern different activities, tools and organizations: in the first, AI was inside a protocol carefully designed for images and readings; in the second, professionals had a chatbot available on simulated cases. They teach the same thing as chapter 24: having the tool does not guarantee that it is integrated into the work.

A Queue Does Not Disappear by Itself

Take any service: a hundred requests arrive each week, eighty are completed. Starting from zero, twenty remain waiting. AI brings capacity to ninety-six. Four remain. The service has improved, but the queue keeps growing.

$$Q_{t+1} = max(0, Q_t + A_t - S_t)$$

Q is the requests waiting, A those arriving, S those completed in the period. The maximum with zero prevents a negative queue, and completed requests cannot exceed those available. If at the start the list was fifty, with a hundred arrivals and ninety-six completions after one week it is fifty-four. The improvement over eighty is not enough to bring it down: to shorten it you must complete more requests than arrive, while maintaining an adequate service.

Even the number of completed requests does not say whether people understood what they were told: they are two different checks. Time saved can become less fatigue for the doctor, more attention for the patient, more capacity for the service; it can also be absorbed by other paperwork. And part of the time spent with the patient is not waste: it is the moment when the doctor listens and discovers what the request did not say. Eliminating it speeds up the procedure and impoverishes care. Distinguishing avoidable waiting from the time quality needs is what really improving the organization means.

Table 55 – Observation outline for distinguishing access, activity and outcome.

Step	Proposed measure	Question to keep
Access	Arrivals and drop-outs	Who does not reach the service?
Activity	Services and times	What is completed?
Quality	Errors and checks	Does the path maintain quality?

Step	Proposed measure	Question to keep
Outcome	Results at follow-up	What changes for people?

And those who check an automated result use skills and time, and may meet a doubt that requires another assessment. Between 1985 and 1987 a radiotherapy machine, the Therac-25, delivered massive overdoses of radiation to some patients in North America, seriously injuring them and killing some, because of a software error. Operators saw incomprehensible error messages, and the system allowed them to go on. Since then that case has been studied in every course on system safety. The lesson still holds: a formal check, without the possibility of understanding, correcting or stopping the process, is not supervision. It is a signature.

The Benefit the Output Does Not Summarize

The model's parameter ω explores the well-being that consumption does not capture. A better service can count without increasing spending; a household can spend more on a necessity without being better off.

A visit near home saves hours of travel, and no price records it; more difficult access can instead take value away from a very fast booking. And a technical possibility can require investment, training and maintenance, and bring to light needs that were previously invisible: greater effectiveness and resource requirements can grow together.

Then there is a way of reading tests that holds for medicine and for every system that flags something. An example, separate from the clinical studies: a thousand people, a hundred with a certain condition, nine hundred without. The test flags eighty of the hundred cases and, by mistake, eighteen healthy people. That leaves

twenty unflagged cases and eight hundred and eighty-two correct negatives.

Table 56 – Example of a test. The condition and the numbers are for teaching; the table does not describe a real disease or the performance of a system usable for clinical decisions.

Flag	Condition present	Condition absent	Total
Positive	80	18	98
Negative	20	882	902
Total	100	900	1,000

Sensitivity is the share of true cases the test finds: eighty out of a hundred, 80 percent. Specificity is the share of healthy people the test does not flag: eight hundred and eighty-two out of nine hundred, 98 percent. The two percentages start from different groups. And 98 percent specificity concerns healthy people, not the probability that a flagged person is ill: changing the starting group changes the question.

Among the 98 people flagged, how many are really ill? Eighty out of ninety-eight.

$$PPV = \frac{80}{80 + 18} \approx 81.6\% \ (example)$$

About 81.6 percent. The denominator defines the question. If the cases become ten in a thousand, with the same sensitivity and the same specificity, you can expect about eight true flags and about twenty false ones. Among those flagged, the ill become fewer than a third. It is the same mechanism as in chapter 17: when what you are looking for is rare, even a good test produces many false alarms. And the twenty cases the test did not find remain.

Go back to the study on doctors: a difference of two points, with an interval from minus 4 to plus 8. Not even the sign of the

effect is known, and the central value does not make the benefit certain. "Not significant" does not mean "identical": it means the study cannot distinguish between the plausible differences. With fifty doctors, the tool showed no measurable advantage.

Finally, those left out of the measure. Looking only at those who used a service introduces a selection: it describes their experience, not that of all those who would need it. Those who measure the success of a digital booking by counting only those who completed it see a success, and do not see those who tried and gave up, those who did not know how to use it, those who did not even try. A service open to all on paper is not a service reachable by all.

Imagine a desk that answers in a few seconds and finds an appointment at once. If the available visits stay the same, administrative time falls, and the real limit, how many visits there are, stays where it was. If the time freed allows staff to take better care of patients, there is an additional benefit. And for those who use the service, the possibility of correcting a booking and a channel other than the screen also count. Booking is not being treated.

Every service needs equipment, infrastructure and energy. The next chapter counts how much, keeping separate what is technically needed and what is actually obtained.

Chapter 29

The Energy Use of AI

Artificial intelligence seems immaterial: an answer appearing on a screen, a service "in the cloud". In reality every answer passes through buildings full of processors, cooled day and night, connected to an electricity grid. Every digital service depends on physical machines and infrastructure, and even when it reduces work in one phase it continues to require energy and capacity, in the place and at the time it is used. To evaluate its weight, three things must be distinguished: power, total consumption, the result obtained.

The Moment and the Year

Power, measured in kilowatts or megawatts, is the rate at which energy is used. Energy, measured in kilowatt-hours or terawatt-hours, is the quantity accumulated over time. Ten megawatts and ten megawatt-hours are not two ways of saying the same thing. Think of a tap: power is how far it is open, energy is how much water you have collected, which also depends on how long you leave it open.

$$E = P \times t$$

E is energy, P power, t time. If power changes over time, the consumption of the different intervals must be added up. For example, a plant running at a constant ten megawatts for all 8,760 hours of the year consumes 87,600 megawatt-hours, that is, 87.6 gigawatt-hours. But the value written on the nameplate can indicate capacity, peak or actual draw: a plant does not necessarily work at maximum all year. And an annual total does not describe peaks: two plants with the same annual consumption can require different grids, if one consumes everything in a few hours and the other steadily.

Table 57 – Units and boundaries of the energy account. The definitions help read the sources; the numerical example does not calibrate the economic model.

Measure	Unit	What it describes
Power	kW, MW	Rate; capacity and actual use distinct
Energy	kWh, MWh, GWh, TWh	Quantity over the period
Conversions	1 MWh = 1,000 kWh	1 GWh = 1,000 MWh; 1 TWh = 1,000 GWh
Emissions	kg or tonnes	They require a factor consistent with the account

How much energy do the world's data centers consume today? In the report Key Questions on Energy and AI of 16 April 2026, the International Energy Agency estimates 485 terawatt-hours in 2025, and projects about 950 in 2030. Almost double in five years. It is the update of the estimate cited in chapter 14, which started from 415 terawatt-hours in 2024 and reached 945 in 2030: the starting year and the date of the projection change. To give an order of magnitude, the Agency itself compares the consumption projected for 2030 to the current consumption of the whole of Japan. The numbers concern all data centers, not only AI; but according to the Agency, AI is the main driver of growth.

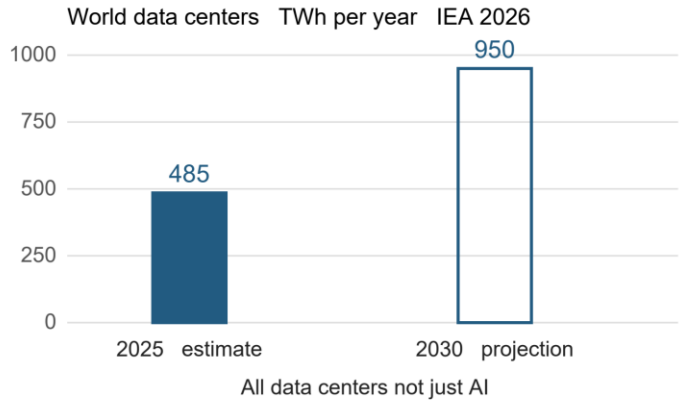


Figure 46 – IEA 2026: world electricity consumption of data centers. Solid bar: historical estimate for 2025; hollow bar: central projection for 2030. Terawatt-hours per year.

The 2025 figure is an estimate of the past; the 2030 one is the Agency's projection. The consumption of a local plant is not obtained by dividing the world total by the number of buildings: you need the plant's data and the load over time.

"How much energy does an AI answer consume?" It is a question that seems simple and is not. You have to define the system, the length of the answer, the result, the data, the discarded attempts: a short question and a long task do not weigh the same. You have to separate the training of a model from its daily use, and decide whether to count only the processors or also cooling and the plant's other services. The book does not assign a universal consumption to every answer, because there is none. An example with declared units, by contrast, makes it possible to understand how the efficiency of a single activity combines with total volume.

Less for Each, More in Total

Here is the chapter's most important mechanism. A service consumes one unit of energy for each of a hundred activities: total, a hundred. A more efficient technique arrives, and consumption per activity is halved. If the activities stay at a hundred, the total falls to fifty. But if, since it costs less, two hundred and fifty are done, the total rises to a hundred and twenty-five.

$$E = e \times N$$

e is the energy per activity, N the number of activities, all considered the same in the example; if the types of activity change, an account for each type is needed.

Table 58 – Example of efficiency and quantity. The energy units are conventional. The growth in activities is assigned, not estimated by IEA research or by the model's software.

Case	Energy per activity	Activities	Total
Before	1	100	100
Efficiency, same quantity	0.5	100	50
Efficiency, more activities	0.5	250	125

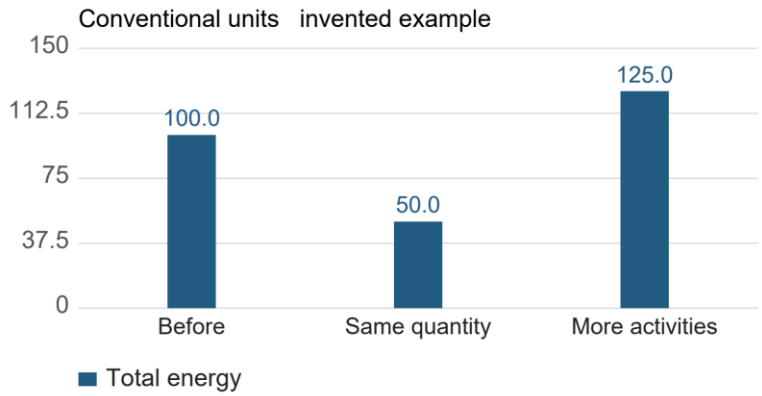


Figure 47 – Totals from Table 58. Improving the efficiency of each activity does not guarantee a decrease in total consumption. Teaching example, not a forecast of future AI use.

More efficiency, more consumption. It is not a curiosity. In 1865 the English economist William Stanley Jevons, in his book on the coal question, observed that more efficient steam engines had not reduced coal consumption in Britain: they had increased it, because using steam had become worthwhile for many more things. Since then people have spoken of the Jevons paradox, or the rebound effect.

The logic is the same as for work: if each task requires less energy, or fewer hours, the total also depends on how much is produced. A lower cost makes possible uses that were previously too expensive, and that is what happened with steam.

The Wire Arrives Somewhere

Energy must arrive where it is needed, when it is needed. Two identical plants can find a ready grid in one place and a saturated grid in another: you need connections and capacity where the demand is, not just energy available elsewhere. And a finished building is not a working service: connection to the grid can require further work, and years. These timescales can slow AI adoption.

And work? A construction site can require many people during construction and few during operation. A figure limited to construction can suggest jobs that do not last; a figure limited to operation forgets work that was really done. Period, duration and skills are part of the account. Those who announce how many jobs a data center will create must say whether they are counting the construction site or the years afterwards.

Then there are emissions. The simplest account multiplies energy by a factor that depends on the place, the period and how the electricity is produced.

$$M = E \times f$$

M is the mass of emissions, E the energy, f the factor. For example, a hundred thousand kilowatt-hours with a factor of 0.1 kilograms per kilowatt-hour make ten thousand kilograms, ten tonnes. With a factor of 0.6, sixty tonnes. Same energy, six times the emissions.

And the operating account is not everything. Constructing the buildings, producing the processors, disposing of them, water and other resources require their own data and boundaries, and a result on one part does not zero the others.

Who pays for the connection? A new power line can serve mainly a data center, or a whole network of users. If the cost ends up in everyone's bills and the benefit goes to a single company, it is a transfer; to discuss its effects you have to follow costs and benefits over the whole life of the work.

In the book's model, energy costs are treated separately, as declared in chapter 6.

Energy in a Firm's Budget

If producing something costs a hundred, of which ten is energy and ninety other items, doubling the cost of energy alone brings the total to a hundred and ten: plus 10 percent, not plus 100. If meanwhile AI saves eight on another item, the total is a hundred and two; if it saves fifteen, ninety-five.

And consumption and spending are not the same thing. You can consume more and spend less, if the price falls enough; or spend more while consuming the same. Terawatt-hours measure physical quantities; the bill also depends on prices and contracts.

The hour counts too. If the emission factor changes during the day, the total is calculated by adding up the energy and factor of each interval; a simple average can weigh badly. Another example: a thousand kilowatt-hours in two periods, one with a factor of 0.1 kilograms per kilowatt-hour, the other 0.6. With five hundred kilowatt-hours per period, emissions are three hundred and fifty kilograms. With nine hundred in the first and a hundred in the second, a hundred and fifty. Same energy, less than half the emissions. Place and hour enter the account. Some computing jobs can be postponed, others require continuity.

The right question, therefore, is not "how much energy does AI consume", but which system obtains which result, with which resources, compared with which alternative. Looking only at the extra consumption hides benefits and savings; looking only at the advantages hides the costs.

Take a library that used to translate a few notices and now, with a more efficient service, also translates catalogs, guides and teaching materials. Each translation consumes less energy than before, but many more are done. It is Jevons's rebound effect in a library: an improvement in the single operation does not decide total consumption on its own.

Power and consumption, energy and emissions, construction and operation describe different things. And all happen in a place. Territory too changes the possibility of using a technology, and of reaching a job. The next chapter links infrastructure and people's days: the same job can be reachable in one place and unreachable in another.

Chapter 30

Mobility and Territorial Differences

In 1968 the American economist John Kain published a study that gave a name to a problem many saw without being able to state it. In the postwar years the factories and offices of many American cities had moved to the suburbs, while many workers, especially Black workers, stayed in central neighborhoods, without cars and with public transport that did not go where the work was. The jobs existed. The people existed. Between them was distance. Later that idea took the name of spatial mismatch.

A job forty kilometers away, for someone with a car, a nearby nursery and nobody to care for, is an opportunity. For someone who depends on the bus and has to pick up a child at four, that job does not exist. Distance on the map does not say how much it costs, how much time it takes and how hard it is to cross. An opportunity can exist in a region without being compatible with a person's day. To read the effects on territories you have to go inside the model's large averages, and distinguish income, services and the possibility of reaching them.

The Region Is Not a Person

The model's three regions have different adoption timing, income distributions and public resources. Within the same large region there are territories with very different connections, firms, schools and housing. For a neighborhood you have to look at where jobs arise, who manages to reach them, who receives the income from ownership, and which services help those who have to start again. Even a change of location or shift can make impossible a journey that used to work: the same distance weighs differently depending on each person's timetable, transport and responsibilities.

Imagine two cities, A and B. In both, a household earns 1,800 euros net a month for the same hours of work. In city A housing

costs 600 euros, transport 100, the childcare needed to be able to work 200. In city B, 900, 200 and 250.

Table 59 – Two monthly budgets. Same net income, different costs. Conventional monetary units; no real city or household.

Item	City A	City B
Net income	1,800	1,800
Housing	600	900
Transport	100	200
Care	200	250
Residual	900	450

$$R = N - A - T - C\ (example)$$

N is net income, A housing, T transport, C the care needed to make that job possible; R is what remains. In city A 900 euros remain, in B 450. Same salary, half the money available. And the remainder does not measure well-being: with that money you still have to eat, pay for energy, face the unexpected.

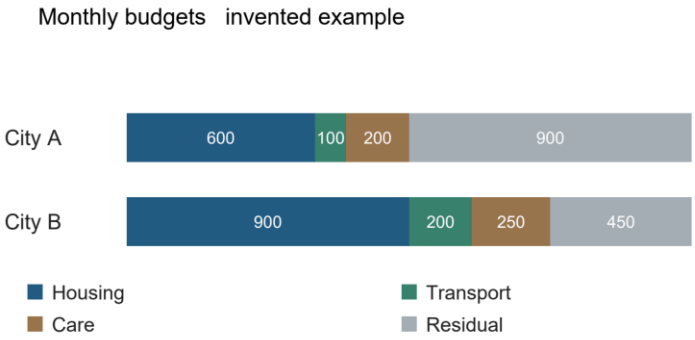


Figure 48 – Breakdown of the two monthly incomes in Table 59. Teaching example; the residual coincides neither with saving nor with total consumption.

City B might offer something else: career prospects, better services, other advantages. And the quality of housing and the safety of the journey count beyond their price too.

Then there are the minutes that do not enter the wage. Two people work eight hours and earn a hundred euros a day. The first travels half an hour in all, the second two hours.

$$q = \frac{N}{h_l + h_v}$$

N is daily net income, h_l hours of work, h_v hours of travel, q the pay for each hour committed. The first person receives about 11.76 euros for each hour of their working day, the second 10. The same contract, two different pays. It is not the contractual hourly wage, because it includes time that is usually not paid; and it is not the whole day, because preparation, waiting, care and rest are missing.

An offer is a set of conditions. A new job can increase a firm's demand without meeting anyone available at that hour, in that place, with those constraints.

A Country Inside the Numbers

What happens in Italy? The annual Eurostat series kept in the project make it possible to follow employment and participation between 2022 and 2025 in four age groups.

The dataset is called lfsa_pganws and counts the population in private households by citizenship and labor status. The processing uses the total of men and women and of all citizenships, in thousands. For each group it first adds up the five-year classes and then divides the employed by the aggregated population: a simple average of rates would give too much weight to the smaller classes. Participation instead uses the active population in the numerator.

Table 60 – Observed changes 2022–2025 in percentage points, from the Eurostat copy archived on October 10, 2026. Italy and EU27_2020; the EU aggregate includes Italy. Employment and participation have the population of the age group in private households as their denominator.

Age	Empl. Italy Δ pp	Empl. EU Δ pp	Partic. Italy Δ pp
15–24	−1.88	−0.18	−3.40
25–34	+2.41	+0.59	+0.67
35–49	+2.31	+0.82	+1.18
50–64	+4.98	+3.13	+4.17

Between 2022 and 2025, in Italy the employment rate of 50–64-year-olds rises by 4.98 points, that of 25–34-year-olds by 2.41 and that of 35–49-year-olds by 2.31. That of young people between 15 and 24 falls by 1.88 points. In Europe, for the same group, it falls by 0.18. In the last year, from 2024 to 2025, young Italians lose 1.78 points and young Europeans 0.38. Weighting each year as in chapter 3, young Italians lose 1.35 points a year, young Europeans 0.32.

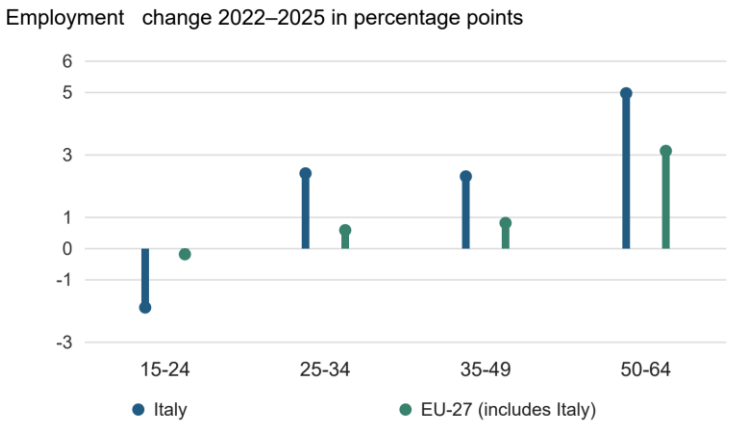


Figure 49 – Change in the employment rate between 2022 and 2025, by age. Italy and EU-27, percentage points. Processing of the annual lfsa_pganws counts; same scope in the two geographies.

The total improves. Young people lose. It is the same pattern that American data show for jobs exposed to AI, and the one the German study on robots found. The labor force participation of young Italians falls by 3.40 points from 2022 to 2025, 1.91 a year with the weights: it is not only that they find less work, some stop looking for it.

With the weights, in Italy the fall for young people is more than four times the European one. The growth in older employment does not offset a more difficult entry: the two measures concern different people. These data say that in Italy the entrance door for young people was already narrowing while the others grew.

The comparison starts from 2022, after the 2021 change of method flagged in the series, and weights the years as in chapter 3. The observations chosen have no break or low-reliability flags. Source and metadata: Eurostat, https://ec.europa.eu/eurostat/cache/metadata/en/lfsa_esms.htm.

A second Eurostat archive follows the moves between employment, unemployment and inactivity from one quarter to the

next. The imported file covers Italy and the European Union in 2024–2025, but without ages: it does not allow checking the entry of young people or distinguishing first job and return. For that question Istat makes available to researchers, with a dedicated request, labor force files three months apart. The protocol in chapter 18 indicates what to link: age, occupation, education and territory, together with AI exposure. Source: https://www.istat.it/microdati/rilevazione-sulle-forze-di-lavoro-dati-longitudinali-a-3-mesi-di-distanza/.

A national statistic indicates where to look for a problem. Routes, stops and timetables explain how a course open to all can remain hard to reach.

Two Precedents and a City That Changes Residents

Two studies on industrial robots help to understand what happens to a territory when automation arrives. They concern a technology different from generative AI, but the mechanism is the same.

Table 61 – Two studies on industrial robots. Summary of the original publications; no estimate of the effects of generative AI and no causal attribution of differences between countries to a single institution.

Study	Scope	Summary result
Acemoglu and Restrepo 2020	Robots, US local markets	Employment and wages falling in the estimate
Dauth and coauthors 2021	Robots, German data	Services offset manufacturing; young entrants hit hardest

In the United States, Acemoglu and Restrepo, in the Journal of Political Economy in 2020, studied local labor markets between 1990 and 2007. One more robot per thousand workers reduced the employment-to-population ratio by 0.2 points and wages by 0.42 percent. Exposure combines the spread of robots across sec-

tors with the industrial structure of each place. In Germany, Dauth, Findeisen, Südekum and Wößner, in the Journal of the European Economic Association in 2021, used administrative data: jobs lost in manufacturing were offset by services, but the cost fell mainly on the young people who were supposed to enter. For those already employed, automation came with more stability in the firm and new tasks.

The German lesson is the most important. The total can recompose across sectors while the entrance door closes, and the job lost and the job created do not belong to the same person.

A neighborhood, moreover, can look richer because its residents have changed, not because those who were there are better off. Ten households with an income of a hundred. If all rise to a hundred and ten, the average grows by 10 percent. If instead two leave and two households with two hundred arrive, while the other eight stay at a hundred, the average rises to a hundred and twenty.

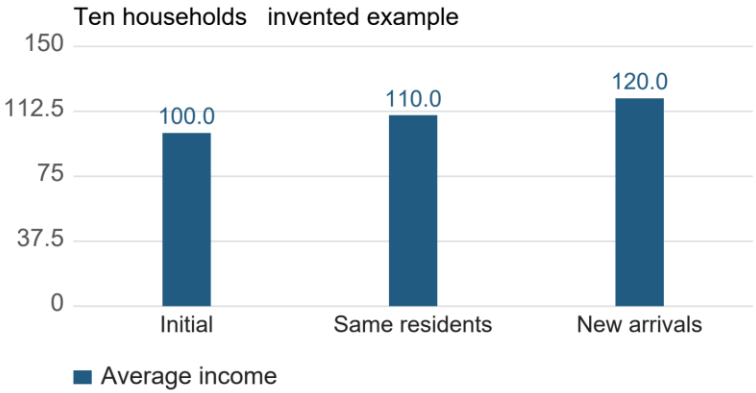


Figure 50 – Average income in three example situations: initial 100, everyone rising to 110, replacement of two households on income 100 with two on income 200, average 120. Ten households in each situation; no measure of a real neighborhood.

The second growth is bigger, and it has not improved the life of anyone who was there. Of the two households that left nothing is

known, if observation stops at the neighborhood boundary. The local measure describes who lives in a place at a moment; the longitudinal one follows the same people over time.

"Let them move where the work is" is easy to say. But the future benefit of moving does not guarantee that it is possible today. Deposits, travel, weeks without pay require money in hand, and caring for children or parents can impose other constraints. The best-known research on place as destiny is Moving to Opportunity: in the United States some poor families received, by lottery, vouchers to move to better neighborhoods. Chetty, Hendren and Katz, in the American Economic Review in 2016, found that children who moved before age thirteen went to college more often and earned more as adults; for adolescents the effect was slightly negative. Age and duration of exposure count. But it concerns children, not adults pushed out of work by AI.

Helping people and helping places are two different strategies. Income support can make searching possible. A transport link widens the reachable jobs. A care service frees hours that otherwise would not exist. An investment can create demand, if the service finds customers and the tasks require people. Each has a cost, and the advantages can be eaten up by higher rents. And a territory can host production and receive little of its income: an access policy does not correct ownership, and an ownership policy does not build the bus. The actions meet, but with different goals.

Two offers with the same salary can leave different money in your pocket if rent, care and transport change, as the budget of the two cities shows. A bus compatible with shifts can widen the reachable jobs without changing the qualification required. A library with space and connectivity helps to send an application, but it does not create the job to apply for. These services stand between the opportunity on paper and the real possibility.

Evaluation must link personal conditions and territory: income, hours, skills and care on one side; firms, jobs, connections and services on the other. And it must observe the encounters and the obstacles, including the people who never reach the counter. The next chapter builds a check of interventions along this sequence.

Chapter 31

Impact Evaluation of Interventions

Every government that launches a training program announces a success. A thousand enrolled, eight hundred graduates, six hundred employed. They look like clear numbers, and they are not. How many of those six hundred would have found work anyway? How many took a job that would have gone to someone else? And the two hundred who dropped out, who were they?

In the United States, at the end of the 1980s, the government did something rare: it evaluated its main training program for the unemployed, the Job Training Partnership Act, with a real experiment, assigning by lottery who could take part and who could not. The results, published in the 1990s, disappointed many expectations: modest gains for adults, and negligible effects for young people who had left school. Without the lottery, counting only the participants who had found work, the program would have looked like a success.

An offer of training is not attendance, attendance is not learning, and learning does not guarantee a job that lasts. A serious evaluation follows the whole chain, and distinguishes those looking for a first job, those returning and those who want to improve a job they already have. In a transition like that of AI, a lot of money will be spent on programs of this kind. This chapter turns the questions into results to measure, denominators and comparisons.

A Result Before the Number

Sustainable hours, sufficient income, recognized skills, an employment relationship that lasts: these are different changes, and indicators must be chosen so as to keep their meaning for people. Research translates them into measures: employment, hours, income, duration. Quality and autonomy require other information,

as chapter 25 showed. The sheet that follows links indicators, observation points and denominators.

Table 62 – Proposed measures for following an intervention. The illustrative horizons must be defined before the evaluation; they are not data from an observed program.

Step	Proposed measure	When to observe
Access	Eligible, offer, take-up and drop-outs	Before and at the start
Learning	Consistent test, with and without assistance	Before and after
Entry	Employment, first entry, return	6 and 12 months
Retention	Duration, exits and new periods without work	12 and 24 months
Conditions	Hours, income and job quality	At every check

Evaluation must also include those who do not even manage to enroll: studying only participants misses the problem of access. A program can work very well for those who take part and reach few of the people it was created for.

The studies on assisted learning in chapter 24 showed that better performance does not always guarantee better independent skill. And even when learning is solid, you have to see whether it meets demand from firms. Katz, Roth, Hendra and Schaberg, in their 2022 article in the Journal of Labor Economics, examined randomized evaluations of programs tied to specific sectors, for low-wage workers. Those programs combine initial screening, occupational and transferable skills, and support services. Participants earned more partly because they entered better-paid jobs, not only because they worked more.

How much more? It depends on the version: 12 to 34 percent in the 2022 article, 11 to 40 in the summary of the 2020 working

paper, 14 to 38 in an IMF summary. They work because they lead towards a sector that is hiring: the same name can indicate different content.

The Missing Comparison

The underlying problem is always the same. You see a person who did the program, but you will never see the same person, in the same period, without the program. Part of the change may depend on the market, on prior skills, on contacts that have nothing to do with the offer. It is the problem of the counterfactual.

Comparing those who take part with those who do not does not solve it. Those who enroll may have more time, resources, motivation, information. Or more difficulties, if the program is aimed at those with fewer opportunities. The cleanest way is to draw lots for who receives the offer, as the American government did.

In an imaginary experiment, two groups of a hundred eligible people are assigned by lottery: one receives the program offer, the other the ordinary services. After a set period, fifty-eight people are working in the first, fifty in the second. The difference is eight points.

Table 63 – Experiment with random assignment of the offer. One hundred people per group; eight percentage points is a teaching difference, with no confidence interval calculated.

Assignment	People	Employed	Rate
Program offer	100	58	58%
Ordinary services	100	50	50%
Difference	–	8	8 points

$$\Delta = 58\% - 50\% = 8\,points$$

Δ is the difference between the means of the two groups, calculated on everyone assigned, including those who never attended. It is the effect of the offer, what statisticians call "intention to treat".

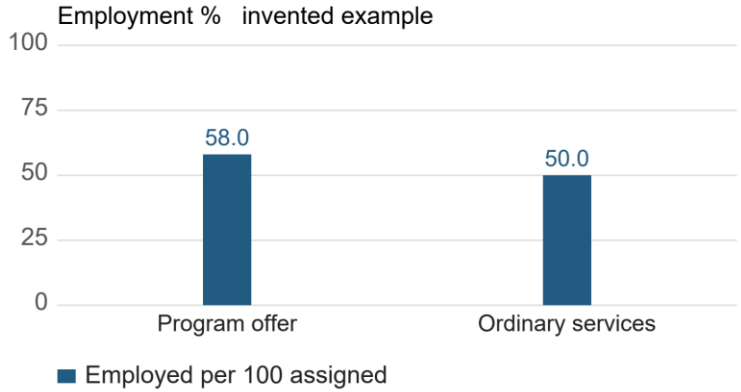


Figure 51 – Employment in the two groups of Table 63. Example, not an individual forecast or a probability of success for a real program.

Here is the point that almost all press releases forget. Fifty people work even without the program. The program cannot take credit for all fifty-eight: it can take credit for eight. And an average effect of eight points does not say which people changed their destiny thanks to the offer.

If only sixty of the hundred people offered the course actually attend, you can divide eight by 0.6 and obtain about 13.3 points. Is it the effect on those who took part? Only if the offer acts solely through participation. A program with interviews, support and meetings with firms can help even those who do not finish the course: attributing every change to the lessons alone would hide the real treatment.

Following people over time is difficult: they move house, change number, stop answering. Of a hundred people assigned, eighty respond and forty-eight of them are working: among respondents, 60 percent are employed. But what about the twenty missing?

$$48\% \leq q_0 \leq 68\% \, (example)$$

q_0 is the share of employed among all hundred. If none of the twenty missing works, it is 48 percent; if all of them work, 68. Twenty points of uncertainty, just because twenty people did not respond. And often those who disappear are precisely those who did worse.

If the Job Only Changes Name

A program can help a person get a job that would have gone to someone else. For the one who gets it, it is a real benefit. For the economy, there is no extra job. It is a problem of displacement: an intervention can change who is hired without changing how many. Another can help firms grow, or fill jobs that would have stayed empty. Yet another can have effects on those who do not take part, through contacts, wages and competition. That is why one person's result does not transfer automatically to a whole city.

The lottery between the two groups must also take this into account. If the offer made to one person changes another's result, the simple comparison may not describe what would happen by extending the program to everyone.

This is the fundamental limit of training in a transition like that of AI. If there are not enough jobs, learning better helps win the race for one of the remaining jobs, but it does not add jobs. Training must be evaluated separately from job creation, and it cannot replace it.

Those who evaluate a program must write beforehand what they will measure, and when and how they will publish the results. That way unfavorable conclusions too become checkable. A protocol written beforehand limits the opportunistic choice of ages, territories, indicators and dates: search through enough compari-

sons and something favorable will always turn up by chance. And if the main result is negative, it must remain visible.

And the cost? Imagine a program that costs 200,000 euros for a hundred people and produces eight more employed people than the comparison group.

$$c = \frac{200{,}000}{8} = 25{,}000\ (example)$$

Twenty-five thousand euros for each additional employed person.

Is twenty-five thousand euros a lot or a little? It depends on how long that job lasts, how much it pays, what other benefits it brings, what the alternatives would cost. A contract of a few days and a job held for years make very different accounts. And if the cost were divided by all fifty-eight employed, instead of by the eight additional ones, the program would look seven times more efficient than it is. It is the most common trick in communicating public policies.

Imagine, finally, a course that ends after the last useful bus. Those without a car stop coming, even before the final test. If only those who reach the end are counted, the evaluation selects the people with fewest obstacles, and presents a success larger than that of the offer. Following everyone to whom the course was offered, as in the example in the table, keeps the boundary of the comparison. Recording the reason for dropping out then helps to correct the organization.

Access, learning, entry, retention and conditions must remain verifiable over time. An offer can create a possibility without having solved the transition. The next chapter changes scale: from single interventions to what the book, as a whole, knows and does not know.

Chapter 32

Results and the Limits of Knowledge

It is time to put the pieces together. Not to reach a single verdict, "AI is good" or "AI is bad", but to say clearly what has emerged, what has not, and what could change our minds.

The Question That Remains

The question of this book is not whether AI can do things. It can, and more and more. The question concerns the work needed to produce, the possibility of entry for those starting out, the return of those who leave, incomes, the resources for services and protection. Technical capability answers only part of this whole.

The model's answer is clear: the value produced by AI returns to households and the state only in part, unequally and late. It is what happens if the rules stay as they are today.

Output grows, the purchasing power of some groups falls, and protecting an income remains different from opening a job. The recovery of work is partial, and follows different timing from region to region. Two dates are enough.

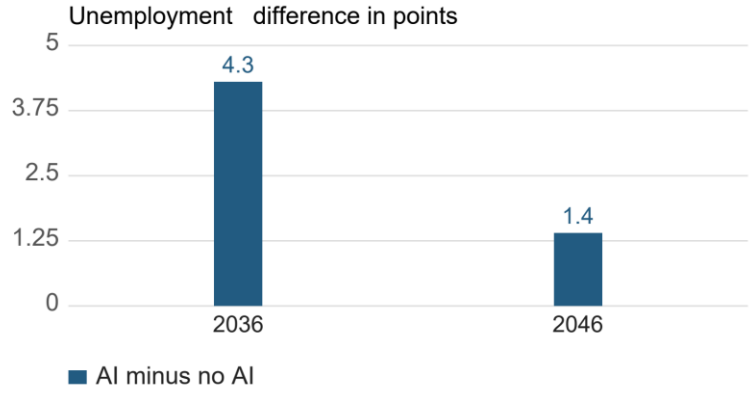


Figure 52 – Median unemployment differential in the advanced economies, AI minus no-AI benchmark: 4.3 points in 2036 and 1.4 in 2046. Model simulations. Two selected years, with no interpolated trajectory and no bands available in this summary.

In 2036, in the advanced economies, unemployment is 4.3 points higher than in the world without AI; in 2046, 1.4. It falls,

but it does not return to zero. A more productive economy can, twenty years later, need less work than the one that did not adopt AI.

What Did Not Emerge, and What Did

One thing, at the start, seemed certain: that falling wages would make consumption collapse, and collapsed consumption would make profits collapse. It does not happen. Demand holds, profits grow. It must be said clearly, because it is an initial hypothesis that the numbers refuted. In the worst of the 4,000 simulated futures demand falls short by 2.5 percent of output, and in seven futures out of ten it never falls short.

The concern for the public accounts, on the other hand, was right in direction, but secondary to ageing. The model started from pressures on the accounts already present, and AI adds a blow to a system that did not start from a blank page: it changes the bases on which taxes are paid and who receives income. In 2046 the debt of the advanced economies is 11.2 points of GDP higher.

Concentrated losses strain trust in institutions.

Then there are two different promises: ownership and a job. Distributing value through ownership, with dividends and funds, as in chapters 13, 22 and 27, protects income. It does not provide a job.

The second indication concerns timing, the breadth of automation and the way the labor market adjusts. In the analyses of chapter 16, the speed of diffusion and the breadth of substitution explain an important part of the variability of unemployment; wage rigidity, new tasks and firms' reaction to the cost of labor weigh on households' lasting results. And none of the tools tested brings unemployment back to the level of the world without AI. Training

and new tasks come closest; postponing adoption is among the least effective.

The third indication concerns the time to build. The unemployment peak in the model arrives around 2039. Measures, services and administrative capacity take years: they must be prepared beforehand. The package in chapter 22 protects consumption more than it recovers work, and the hardest question remains open: who opens the jobs, and how do you check that they are reachable. Distributing value better can defend income; to get in you also need demand, a match between tasks and people, real opportunities. Chapter 24 showed it: even with the best incentive, first entry of young people remains well below the world without AI.

What Would Change the Conclusion

A checkable conclusion says which new observations could weaken it. Here they are.

A stable increase in first hires in jobs that really use AI, with credible comparisons and acquired skills, would weaken an overly general version of the risk for young people. Growth in new tasks broad enough to offset those lost would force a revision of the parameters. A wider distribution of capital income, observed and lasting, would change the households' account.

But it also holds the other way round. Slower returns, lost hours hidden inside the count of the employed, even more concentrated benefits, services promised and never delivered would make the conclusion harsher.

Table 64 – Proposed questions for updating the analysis. Indicators and comparisons to be defined in a protocol; not validated thresholds or results of future research.

Question	Information to look for
Entry	First hires and verified adoption, with a comparison

Question	Information to look for
Re-employment	Same people: duration, hours, income and quality
New tasks	Additional tasks and jobs relative to substitution
Distribution	Wages, ownership, returns and transfers
Policies	Causal effects, costs, access and duration

Reachable sources, explained formulas, consistent definitions: that is how a conclusion can be criticized.

A useful forecast leaves visible the point at which new evidence could modify it.

Go back to the sign workshop. A faster delivery can raise the margin, lower the price or free hours. The firm's account measures costs and revenues. The work account measures hours and wages. The household account measures the resources really available. If the software fee pays a distant capital, there is a fourth account: where the income goes. These accounts can improve to different degrees, or go in opposite directions. The book's conclusion holds them together.

The next chapters summarize the results, with their numbers: work and incomes, the circle that links them, the policies.

Chapter 33

Summary of Results: Work and Incomes

If you were to remember only one chapter of the results, remember this one. Here is what happens to work and incomes, in the model's numbers.

In 2036 world GDP is 1.3 percent higher than in the world without AI. The share of income going to labor is 3 points lower. In 2046 GDP is 3.2 percent higher, and the labor share 5.9 points lower.

In other words, the pie grows, and labor's slice shrinks. It is not a paradox. It is exactly what happens when machines do a growing part of the work and their income goes to those who own them.

A lower share does not measure each person's wage. A firm that automates and sells more can increase its total wage bill by less than the value produced, or reduce it. To know how households are doing you also have to look at ownership, transfers and prices. That is why the book looks at consumption.

Who Gains and Who Loses

In 2036 the world's poorer half consumes 1 percent less than in the world without AI. The middle forty percent 0.3 percent less. The richest ten percent 3.1 percent more.

The poorer half loses in 92 futures out of a hundred. In 2046 its loss reaches 3.3 percent, and remains negative in 89 futures out of a hundred.

Growth does not close the gap. Not even in twenty years. And these are the numbers of the cautious scenario: the model excludes an AI able to do almost everything.

A result that holds in nine futures out of ten, as all the parameters vary, is not a detail.

The Door Closes Before the Total Moves

For those looking for a first job, the number of employed arrives late. Chapter 8 followed young people between 15 and 24 in the advanced economies, in the scenario in which retraining is slow. In 2036 the share of those who manage to get in is 14.3 points lower than in the world without AI. Measured on the whole available population, it is about 1.37 fewer entries per hundred people every year. Two ways of counting the same phenomenon.

It is the mechanism that the American data of the last year already show for young people in exposed jobs: a transition can close entry-level jobs without firing those who already have experience. Total employment stays almost still. The entrance door narrows.

A library that automates initial cataloguing can keep its experienced staff and remove the tasks that used to be given to beginners. A café that automates orders and payments can remove the first rung from which the rest was learned. The question is: where will the experience needed for later jobs be gained, if the work through which it was gained disappears?

The Time of Loss and the Time of Recovery

In the advanced economies additional unemployment compared with the world without AI is 4.3 points in 2036 and 1.4 in 2046. The peak arrives around 2039, at about 4.9 points, and stays above four points from 2036 to 2041. Six years.

That unemployment falls at the end does not mean that the same people have returned, nor that they have recovered their previous hours, income and conditions. Chapter 25 showed why: a higher hourly wage can hide fewer hours and less income. A shift that ends after the last bus can make a job unreachable.

Output, access to work and income distribution can separate for years. Output grows, the door narrows, the gap between those who have and those who do not widens. Differences between regions change the intensity, not the direction.

What all this means for a specific occupation is in chapter 36, with a table for nineteen occupations.

Chapter 34

Feedback Between Employment and Incomes

A calm result can hide a danger. That is the lesson of this chapter, and perhaps the most important one for those who will have to decide.

Chapter 26 closed the circle between work and incomes: those left without work weigh on wages and benefits, and lower wages weigh on demand. In the reference scenario, closing the circle changes little. The consumption loss of the world's poorer half in 2036 moves from 1.03 to 1.39 percent.

If you stopped here, you would conclude that the circle matters little. It would be a serious mistake.

When Things Go Wrong

In the scenario in which retraining is slow, the loss moves from 1.43 to 3.15 percent. In the scenario in which substitution persists, from 2.64 to 7.24. In both cases the loss more than doubles.

The circle between work and incomes is almost invisible when the transition works. It becomes devastating when it jams. Those who lose their jobs spend less, firms sell less and hire less, and those looking for work stay out longer. The damage feeds itself.

It is not an abstract hypothesis. After the 2008 financial crisis, in the United States, long-term unemployment exploded: in 2010 almost half of unemployed Americans had been looking for work for more than six months, a share not seen since the postwar period. The longer the wait, the harder it became to get out: employers looked with suspicion at résumés stalled for months, skills aged, savings ran out. The model's difficult scenarios show that, when the return jams, the damage does not add up: it multiplies.

The model says something precise: the future in which things go badly is much worse than the average scenario suggests.

Here is the consequence. Support that is enough when people return to work quickly may not be enough at all when the loss per-

sists. A policy must be tested in the difficult scenarios, not in the one in which the problem seems manageable. It is like a levee: you do not build it for the average flood.

The experiment shows that the damage is amplified.

On to the Policies

The policies in chapter 24 change the opportunities for work. Their benefits are not subtracted from the losses in this chapter: they are two different experiments.

The next chapter compares the policies. Take this question with you: does the measure you choose still protect when the transition goes badly?

Chapter 35

Summary of Policies and Open Choices

Same budget, three tools, three different results. Choosing a policy means first of all choosing what you want to achieve.

The Same Cost, Different Goals

The equal-cost comparison in chapter 24 puts three tools on the same footing: public jobs, an incentive to hire young people, reduced working hours. The budget is 0.7 percent of each region's output, every year from 2027. The figures concern the advanced economies in 2036, in the reference scenario, and measure the effect of the policy in the world with AI compared with the same world without intervention. They are not effects of AI compared with the world without AI.

If you want to get young people in, the incentive wins. First entry of 15–24-year-olds rises by 3.54 points, against 1.93 for reduced hours and 1.28 for public jobs. For employment of 25–34-year-olds the improvements are 2.93, 1.79 and 1.09 points.

If instead you want to reduce the total number of people without work, including those who have stopped looking, the ranking is reversed. Reduced hours lower it by 1.82 points, public jobs by 1.16, the incentive by 0.97. Spending the same amount does not give the same results. Inside the choice criterion lies a priority: reaching those who enter better, or reducing the gap more for everyone. The simulation makes this trade-off visible; the choice remains with those who decide.

The Improvement and the Gap That Remains

The incentive's benefit does not cancel the technological problem. Even if the world without AI received the same incentive, first entry in the world with AI would remain about 11.88 points lower. It is a different comparison from the plus 3.54 of the policy com-

pared with doing nothing. A measure can help and leave an enormous gap: both results are part of the evaluation.

Each tool has its limit, and examples help to see it. A café that receives a subsidy to hire can open one more job, or collect help for a hire it would have made anyway: it is the difference between additional hires and those the model assumes would have happened anyway. A public service can offer a useful job, but the salary paid does not by itself provide premises, equipment and organization. Reduced hours can divide work among more people when shifts and skills allow it.

And the cost for those already working must be kept visible. In the advanced economies, in 2036, hours fall on average by about 2.35 percent; after the assumed public top-up, labor income falls by about 0.94. More people work, and each earns a little less. Germany, during the pandemic, made massive use of short-time work, and jobs held; but it was a temporary crisis, and short-time work served to preserve a job destined to come back. Permanent substitution is something else.

Protecting Consumption, Rebuilding Work

The policies in chapters 13 and 22 also answer a question about distribution. In that different experiment, training and new tasks reduce the rise in advanced unemployment in 2036 from 4.44 to 3.30 points. The full package brings consumption of the poorer half to plus 3 percent in 2036 and plus 3.3 in 2046, compared with the world without AI with the same policies.

The reference matters because part of the interventions redistributes income even without AI. Comparing the world with AI and the policy with the one without AI and without the policy would mix the technological advantage with that of general redistribution. The book presents both comparisons in the respective

chapters; this summary uses the one with the same policy when it asks what remains of the damage attributed to AI.

In the equal-cost comparison the budget builds additional jobs; the ranking holds for the declared goals.

For a library, a workshop or a transport service the benefit can take different forms: longer opening hours, a new skill, a first contract, a lower expense. An intervention that works on one of these dimensions does not necessarily improve the others.

The book's results thus lead to a precise method. When the question is how to use resources, tools are compared at equal cost. When the question is how much of AI's damage remains, the world with and without AI is compared with the same policy. When the goal is to protect households, consumption is what you look at. It is on these comparisons, and on the checks to come, that a decision can have a checkable basis.

So far the model. The next chapter takes the numbers into the occupations.

Chapter 36

The Last Forecast

Up to here the numbers have spoken. From this page on, the forecast speaks.

It is a distinction that must be kept clear, because an honest book does not mix the two. The model showed what happens in a cautious scenario: a powerful but not omnipotent AI, exposure growing by at most 60 percent, no robot able to do almost everything. Even so, the economy grows, the poorer half loses, the door for young people narrows, unemployment rises for a decade and does not return to its starting point.

This chapter goes further. It says what to expect if today's trends continue, and it says so occupation by occupation too, with a table built on published data and rules listed only once. Those who wait for certainty arrive late.

The Forecast

Between 2026 and 2036 artificial intelligence will reduce the demand for human labor in an important part of the cognitive activities that can be standardized. In many areas it will not just be employment growing more slowly: there will be fewer jobs, in absolute terms. New jobs and increased demand will offset part of the losses. They will not reabsorb them all, and not within a time that a life can wait.

In the most exposed economies the consequence will be a long transition: harder entry, wages under pressure in many activities, value ever more concentrated in a few hands.

This forecast assumes three things: that the capabilities of the systems continue to improve, that adoption spreads, that institutions do not change fast enough to compensate. It also assumes a basic, declared choice: that incomplete reabsorption is the scenario to prepare for. Assuming that the market will rebuild on its own

an equivalent job for every job lost would be, today, a bet without foundation.

Much administrative, document-based and digital work will need fewer people. The most vulnerable activities are those in which the result can be produced and checked with information that is already digital: paperwork, standard customer support, documents, ordinary translations, repetitive content, some analyses, some phases of software development.

And AI does not have to be able to do an entire job to reduce staff. Take any office with ten people preparing paperwork and answering requests. If automation allows seven people to handle the same work, the job continues to exist. And three positions disappear. If volume grows enough, those positions can remain. The real question is not whether at least one human task remains. It is how much work grows after each person has become more productive.

The contraction will begin silently, with the jobs that are not opened. A retirement not replaced, a contract not renewed, a recruitment canceled. No mass layoffs, no headlines. That is why first entry of young people is the signal to watch before all others. An economy can look stable to those who already have a job and be ever more closed to those looking for one.

There is one thing this book does not do: it does not say how many jobs will disappear in the world. The total balance also depends on population, growth, services and public choices, and no model today can calculate it credibly. Whoever gives that number is guessing. The forecast gives the direction, the mechanism and the share of work each occupation loses, not a global figure.

The signals are already there. In the ADP payroll data collected by Brynjolfsson, Chandar and Chen, updated to August 2026, young people aged 22 to 25 employed in occupations exposed to AI in the United States are 19 percent fewer than they would have

been following the trend of their peers in less exposed jobs. Weighting the years as in chapter 3, they lose 4.2 percent a year against 0.3 for their less exposed peers; in the last year, 4.4 against 2. And the decline is moving up the age range: between 26 and 30, the most exposed occupations lose 3.1 percent in the last year, while the others stay flat. The difference arises mainly from fewer hires: the authors do not yet find widespread layoffs. [36.1]

Massenkoff and McCrory, in March 2026, distinguished what AI could do from what it actually does, also using Claude usage data. They do not find a systematic increase in unemployment in the most exposed occupations, but they find signs of a slowdown in the entry of young people. And they show that real use still covers only part of what is technically possible. It is research produced by a company in the sector, and must be read knowing that. [36.2]

The ILO estimates in 2025 that about a quarter of world employment is to some degree exposed to generative AI, and considers it more likely that jobs will be transformed rather than disappear. [36.3] But a transformed job can require fewer people. That an occupation remains on the list of jobs does not mean it keeps the same number of employed, the same entry route, the same pay.

And capabilities are growing while adoption advances. METR, in its Time Horizon 1.1 update, measures the length of software and research tasks that AI agents can complete with 50 percent success: since 2023 it has doubled about every 131 days. It concerns bounded tasks, but it says that reasoning as if today's machines were tomorrow's is a very optimistic assumption. [36.4]

Together, these data justify forecasting growing pressure on exposed jobs. The further step, that the gap between what AI can do and what it already does will be closed, and that new jobs will not be enough, is the responsibility this chapter takes on.

Why the Tractor Comparison Is Not Enough

The tractor is always cited. It took millions of people off the fields, and those people found work in factories. Bustos and co-authors documented this shift in Brazil: when agriculture became more productive, workers moved towards manufacturing, especially towards low-research-intensity industries, with largely unskilled workers. [36.5] It is a case of movement between sectors. It does not prove that every technology creates equivalent jobs.

And it worked because there was somewhere to go. Those who left the fields found the factory. Those who left the factory found the office. Brain work was for two centuries the refuge: the part that machines could not do.

AI enters precisely the refuge. Those who lose an office job and look for another job in services may find the same automation there. The boundary between the activities replaced and those that should absorb workers moves while people try to cross it. A course that today prepares for a task may find itself, at the end, facing a demand that no longer exists.

And AI does not need a new machine in every workplace. Software is distributed through the network to millions of firms at once, small ones too. The tractor had to be bought field by field. AI comes by subscription.

Software development is the closest case. This book forecasts that within about ten years writing code will stop being a job in its own right: AI tools will do it, and in its place will remain a figure closer to an architect, who decides what to build, breaks the problem down, guides the tools and checks the result. That figure will not return to today's numbers. If one person with agents does the work that today requires many programmers, getting back to the same employment would take far more software, and nobody has shown that demand can grow that much. The first signs are al-

ready there: in the last year, among software developers in the ADP sample, those aged 22 to 34 fell by 4 to 5 percent, while those aged 41 to 49 grew by 2.6; with the weights of chapter 3, 22- to 25-year-olds lose 2.7 percent a year and 41- to 49-year-olds grow by 5.1. Firms keep those who can design and take on fewer of those who write code.

The calculation is the one in Table 65, further on. AI touches 84 percent of a developer's tasks, according to the expert judgment collected by Eloundou and co-authors. Of the exposed tasks, 85 percent are substitutable, and 85 percent of these are automated by 2046: about 61 percent of the tasks pass to machines. The new roles inside the occupation give back part of them, and demand for software, which the US Bureau of Labor Statistics sees growing by 10 percent between 2025 and 2035, another part: jobs fall by 37 percent in 2036 and by 41 percent in 2046. And a problem opens that the market does not solve on its own: tomorrow's architects are today's programmers. If the first rung disappears, where will they come from?

It would have to be revised if demand for software grew enough to reabsorb developers, if firms that really use agents went back to hiring young people, or if agents remained unreliable on complex systems throughout the decade. Diffusion matters too: in 2025 in Europe, 17 percent of small firms used AI against 55 percent of large ones, and old systems, regulated sectors and public offices will keep hand-written code alive longer.

Occupation by Occupation

Manual work remains. For many it is the last safety net: if the office empties, there will always be the plumber, the warehouse worker, the nurse, the bricklayer.

This book forecasts that this refuge too will shrink. The main limit of robots was never the mechanical arm: it was the brain. A robot could repeat a movement in an orderly factory, not deal with a building site, a home, a warehouse where every day is different. Artificial intelligence is exactly what was missing. When a machine that can see, understand and plan is put inside a body that can move, the physical work that seemed protected begins to be protected no longer.

Today's numbers say where we start from. According to the International Federation of Robotics, in 2024 there were 4.66 million industrial robots in operation in the world, 9 percent more than the previous year, almost all inside orderly factories, repeating precise movements. Service robots and humanoids are still at the beginning, expensive and not widespread. But the history of technologies teaches that the initial phase, expensive and clumsy, is also the one in which it is easiest to underestimate them. Korinek and Suh, in their scenarios, show what would happen if automation came to cover almost all cognitive and physical activities: wages could rise as long as something remains that only people can do, and collapse when that something runs out. This book's model excludes that scenario out of caution. This chapter does not take it as certain, but refuses to take it as impossible.

It will not all happen at once or everywhere. Costs, reliability, safety and regulation will slow things down. Many care and relational jobs will hold out longer. But the direction is clear. If AI closes the refuge of brain work and AI with robots begins to close that of manual work, the question that has accompanied every technological revolution, where will people go, is left without an automatic answer.

Table 65 turns the forecast into numbers for nineteen occupations, in the advanced economies: how many jobs there will be in 2036 and 2046 compared with 2026, with AI, robots and the de-

mand for each occupation. There are five rules, and they apply to every occupation.

The first is exposure. For AI it is the share of an occupation's tasks whose time a GPT-type system, with its tools, can at least halve, according to the judgment of human experts collected by Eloundou, Manning, Mishkin and Rock in Science in 2024. For routine office jobs the higher of this and the probability of automation estimated by Frey and Osborne in 2013 counts, which for the office clerk is 0.96. For robots the Frey and Osborne probability for manual jobs counts.

The second is the type of work, which decides how much of the exposed tasks is substitutable and how much of that is actually automated by 2046: 85 and 85 percent in routine office and content work, 70 and 70 in design with AI tools, 55 and 47.5 where responsibility and human contact count, 55 and 70 in manual work, for the part done by AI. For robots the Frey and Osborne probability is already the automatable share of the occupation, and 70 percent of it is automated by 2046. The distinction follows the International Monetary Fund's distinction between exposure and complementarity.

The third is time. Adoption follows the book's S-curve, with central year 2033. For call center agents and office clerks the central year is calculated so that the first year reproduces the decline of customer service in the ADP data, weighted by year as in chapter 3 and measured against peers in less exposed jobs: minus 1.9 percent a year. For software developers the decline concerns only those under thirty, while above thirty employment grows: the central year stays 2033. Robots arrive five years after AI.

The fourth is the new roles inside the same occupation: for every ten tasks lost, one and a half come back, three years later. In the emerging economies everything arrives three years later, in the low-income countries six and a half years later.

The fifth is demand. The work machines take over is not the same as jobs lost when an occupation is needed more and more. Each occupation keeps the growth projected by the US Bureau of Labor Statistics between 2025 and 2035, in the projections published on 21 September 2026, at the same pace to 2046. The Bureau attributes to AI only the decline of office and administrative jobs: for the office clerk and the call center agent demand is held flat, so that automation is not counted twice.

Table 65 – Jobs by occupation in the advanced economies in 2036 and 2046, compared with 2026, with AI, robots and demand for the occupation. AI: Eloundou, Manning, Mishkin and Rock (2024) and Frey and Osborne (2013). Robots: Frey and Osborne (2013). Demand: change 2025–2035 from the Bureau of Labor Statistics, flat for office jobs. Data: aimodel, outputs/occupations.json, the same as the simulator.

Occupation	AI	Robots	Demand	2036	2046
Retail clerk	45%	92%	−0.3%	−30%	−68%
Factory worker	18%	97%	+1.1%	−23%	−61%
Office clerk	95%	–	flat	−50%	−58%
Journalist	83%	–	−5.9%	−45%	−57%
Accountant	94%	–	+5.0%	−45%	−53%
Call center agent	86%	–	flat	−45%	−53%
Driver	31%	79%	+3.8%	−21%	−53%
Translator	88%	–	+2.0%	−43%	−52%
Graphic designer	82%	–	−1.7%	−42%	−52%
Warehouse worker	12%	85%	+1.8%	−18%	−52%
Cook	16%	96%	+12.1%	−14%	−51%
Construction worker	5%	88%	+7.3%	−13%	−46%
Software developer	84%	–	+10.2%	−37%	−41%
Solution architect	82%	–	+7.9%	−23%	−23%
Teacher	52%	–	−0.2%	−10%	−12%
Lawyer	80%	–	+4.7%	−11%	−10%
Electrician / plumber	18%	25%	+8.3%	−2%	−7%
Doctor	54%	–	+3.3%	−7%	−6%
Nurse	57%	–	+5.6%	−5%	−3%

These are jobs in that occupation, not unemployment: people can move elsewhere, if they find somewhere to go. At the top are the jobs that robots reach and office work: the retail clerk loses 30 percent of the jobs in 2036 and 68 in 2046, the factory worker 23 and 61, the office clerk 50 and 58, the journalist 45 and 57. Where demand grows, the account changes. The nurse hands 13 percent of the work to machines, mostly documentation and paperwork, but the care needed increases, and in 2046 jobs are only 3 percent fewer: the time freed goes back to patients. The cook loses half of the jobs and the construction worker almost half; the electrician 7 percent. The doctor loses 6 percent, the teacher 12. The simulator published with the book uses the same data and the same rules, and gives the same numbers.

And the work needed to build all this? Semiconductors, data centers, networks, energy and maintenance require people. But this book's forecast is that this work will be much less than what it replaces: fewer people, with different skills, in different places, for different periods. A robot factory does not employ as many workers as the robots it produces replace. If it did, nobody would buy the robots.

The argument for recovery has a logic. Lower costs lower prices, lower prices broaden demand, demand creates new products and new jobs. Those who become more productive receive more orders and hire. These channels exist, and that is why an hour saved is not automatically a job lost.

But for the recovery to be complete, four things are needed together: enough jobs, soon enough, with the right skills, in the right places. That new activities arise proves none of the four. Assuming all of them as the natural outcome of such a broad transformation is an act of faith, not a forecast.

Jobs supervising AI also have a limit. If one person can check the work of ten machines, supervision remains human, but the

number of people falls. Some supervisors will be more qualified and better paid. They will be few.

And training helps to compete for jobs, it does not multiply them. If many people are prepared for few opportunities, training decides who wins, not how many win. Chapter 31 showed it. Proposing retraining as a general answer, without measuring it, means attributing to it a power nobody has yet demonstrated.

Young People and the Experience No Longer Gained

The most worrying part of the forecast concerns the way people become experts. Many entry-level tasks serve at once to produce and to learn: a first analysis, a draft to correct, the most common requests, the errors to observe. If AI removes them, firms will continue to look for experienced people, and offer fewer and fewer occasions to become one. Today's gain prepares tomorrow's void.

A professional practice can produce more documents with the same experts and hire fewer trainees. A library can automate initial cataloguing and keep only those who assess the difficult cases. The expert and the beginner use the same tool, and do not receive the same advantage from it.

The book's age module, in chapter 8, shows the same mechanism in simulation. The American data show it in reality. They point to the same door.

A postponed entry postpones everything else: independence, a home, the plans that require a stable income. A return to a lower-paid job leaves intact the debts incurred before. The model does not quantify these consequences. They are no less real for that.

The pressure will go beyond those directly replaced. If those leaving office work look for work in less automated services,

competition for those jobs increases, and wages suffer. Other sectors are not a bottomless container.

Care, maintenance, education and local services have enormous needs. But a need becomes a paid job only when someone finances it, organizes it and prepares people to do it. A bank clerk does not become a nurse because nurses are needed. The distance between need and paid demand is the place where policy can change the outcome.

The most dangerous combination is visible growth accompanied by personal opportunities that are ever harder to reach. It can strain trust in education as a road to improvement, and in institutions as a pact between contribution and protection. The social risk is that many people experience progress as something that takes away their power and role, while total wealth grows.

Chapters 33–35 are the starting point. Chapter 33 brings together growth, labor share, distribution and entry. Chapter 34 shows how much the circle between work and incomes weighs. Chapter 35 compares the interventions. From there comes the basis of the forecast: more productive capacity can coexist for a long time with fewer opportunities and with a loss for part of households.

The partial recovery of unemployment remains in the accounts. But it is an incomplete recovery: unemployment does not return to zero, the loss of the poorest continues. And that unemployment falls after a peak does not mean that the same people have recovered work, income and lost years.

The difficult scenarios of chapter 34 show that the damage is amplified precisely when getting back in is hard. Preparing only for the reassuring scenario would be unwise.

And what the calculation leaves out can weigh. The model does not reconstruct the experience lost by those who enter late, nor

every social effect of lost income. These absences prevent a curve that rises again from being presented as a problem solved.

The path to 2046 has three phases. Between 2026 and 2030 the transformation will be seen mainly in hiring: jobs not opened, missed replacements, digital work teams reorganized. Some services will grow, others will need fewer people. The gap between those with experience and those trying to gain it can widen a great deal, even without a visible increase in unemployment.

Between 2030 and 2036 the shift will be from help on single tasks to the reorganization of whole sequences of work. If systems are reliable and cheap enough, firms will redesign complete document and digital procedures, and the pressure will rise on intermediate roles too, not only on entry-level ones.

Between 2036 and 2046 an incomplete adjustment is more likely than spontaneous compensation. New activities will absorb work, but lasting differences between territories and groups will remain. And in this phase AI-driven robotics will extend automation to a growing part of physical work: this is where, in Table 65, retail clerks and factory workers lose more than six jobs in ten, drivers, warehouse workers and cooks more than half. But the employment difficulty forecast does not depend on it. Even if no manual work were automated, the transition described would remain hard.

The years mark phases, not appointments. The thread that links them is one: the capacity to produce grows faster than the capacity of institutions to distribute opportunities and incomes. How it ends depends on the decisions taken along the way, not on the technology.

What Would Refute This Forecast

A forecast that cannot be refuted is worth nothing. This one can be.

It would have to be revised if, in firms that really use AI, output growth sustained employment and working hours over time, with hiring of young people at least equal to that of similar firms that do not use it.

It would have to be revised if new tasks opened learning paths and additional jobs on a scale comparable to the losses. An increase in total employment in a period of strong general growth would not be enough to prove it.

It would have to be revised if the people affected really recovered income and job quality, or if the spread of AI remained held back by costs, reliability and organization much more than forecast here. AI can be technically impressive and have a smaller effect on work than expected.

And it would have to be revised, on manual work, if AI-driven robots remained confined to a few controlled environments throughout the next decade.

If instead the transition went well thanks to effective policies, the forecast of the social outcome would change. But it would not prove that reabsorption was automatic. It would prove that choices matter.

Preparing for this forecast means treating the loss of opportunities as a central problem of growth, not as a side effect. Training addresses part of it. Shorter hours, demand for services, support during the transition and wider access to capital income address other parts. No tool, alone, rebuilds the previous balance.

In concrete terms, the forecast points to three urgencies. The first is the entrance door: if young people no longer find the tasks from which the trade was learned, someone must build other ways for them to learn it, and someone must pay for them. The second is ownership: if value moves towards those who own machines and infrastructure, it must be decided now how to broaden that ownership, before it concentrates entirely. The third is time: ser-

vices, funds and rules take years, and the peak of the transition, in the model, arrives towards the end of the 2030s.

When production requires less paid work, distribution becomes the main question. If a growing part of value goes to those who own technologies and infrastructure, access to that value will depend more and more on rights, contracts and institutions, and less and less on hours worked. A society can turn the time saved into services and free time for all. Or it can let it become exclusion from income for many. Technology makes both roads possible. It does not decide who pays.

This is the book's forecast. Artificial intelligence will significantly reduce job opportunities in exposed activities, and with robots it will also approach manual ones. Reabsorption will be slower and more incomplete than the comparisons with past revolutions promise. What awaits us is an economy able to produce more with less human work. If institutions do not change, the value freed by machines will go to those who own them. For it to go also to those who will no longer be called to work in the same forms, it has to be decided. And it has to be decided now.

Chapter 37

Social Risk

Chapter 36 counted the jobs. This chapter counts what happens to the people who lose them, country by country, month after month.

Chapter 17 said something that still holds: no verified formula turns a point of unemployment into a certain number of protests. This chapter does not make one up. It declares a rule, follows it to the end and shows where it leads. It is the same method as the forecast by occupation: rules written down only once, published data, numbers anyone can redo. The same projection runs on the book's website, on the Social Risk page, with the same data and the same numbers as these pages.

The rule

More AI and more robots mean more people replaced. More people replaced means more social tension. And all of this comes first in the developed countries, because AI and robots enter there first and find more work to take.

Where a country is already tense today, the same replacement weighs more. It is the effect of a sounding board: in an unstable country even a little AI finds ground that amplifies it. But the amplification has a limit. It can make the same loss weigh at most half as much again, never enough for a fragile country with little AI to overtake a stable country that loses much more work.

The pace is not even. As long as only AI is at work, tension grows slowly: AI removes tasks, slows hiring, leaves posts empty without firing anyone, as chapter 36 showed. When robots enter, replacement becomes visible, in warehouses, factories, shops and on the roads. From that moment tension grows at full speed.

The ingredients

The first ingredient is the work taken by AI and robots. It comes from the nineteen occupations of Table 65, weighted by the number of people employed, with the same adoption curves as chapter 36. Emerging economies start three years later, low-income countries six and a half years later. And not all the work of a poor country is exposed: the share of employment exposed to AI is 60 percent in the advanced economies, 40 in the emerging ones, 26 in low-income countries, as in chapter 6. Fewer exposed jobs, fewer people replaced.

The second ingredient is AI use in each country. It is the share of the working-age population that uses generative AI, measured quarter by quarter by Microsoft's AI Economy Institute and published by Our World in Data, from June 2025 to March 2026. After the last observation it grows along the book's adoption curve, up to the share of employment exposed in the region. Within a region, the work lost is shared among the countries according to how much AI they use: those who use more replace more people.

The third ingredient is the starting tension. It comes from the World Bank's index of Political Stability and Absence of Violence, Worldwide Governance Indicators, 2024 estimate, brought to a scale from 0 to 100 where higher means less stable. It does not appear in the tables: it serves only as an amplifier. A country with maximum stability multiplies replacement by 1, one with maximum instability by 1.5.

The fourth ingredient is the pace. Until robots enter, tension grows at a quarter of the speed. The entry of robots, in a region, is the moment they take 3 percentage points of the work of the nineteen occupations: March 2033 in the advanced economies, March 2037 in the emerging ones, December 2041 in low-income countries. After that, tension grows at full speed. Month by month only

the weighted growth is added up, so the accumulated delay remains.

The regions are those of the International Monetary Fund. The advanced economies are the forty countries on its list. The low-income countries are those the World Bank classifies as low income, plus those with lower-middle income in sub-Saharan Africa. All the others are emerging economies.

Everything is measured as the increase compared with January 2026. The tension number is the work replaced since that date, in percentage points, weighted by the starting tension and by the pace. It is not a share of people in the streets. It is a measure of how much replacement a country has to absorb and of how fragile the ground it falls on is.

What the projection shows

Table 66 – Entry of robots, work taken by AI and robots and median social tension by region, in 2038 and 2046. Work taken: percentage points of the work of the nineteen occupations of Table 65, scaled by the region's exposed employment, compared with January 2026. Entry of robots: month in which they take 3 points. Data: aimodel, outputs/social_tension.json, the same as the Social Risk page.

Region	Robots	Work 2038	Work 2046	Tension 2038	Tension 2046
Advanced economies	Mar 2033	39	55	37	56
Emerging economies	Mar 2037	17	35	13	35
Low-income countries	Dec 2041	5	19	2	18

Until 2030 almost nothing visible happens. AI spreads fast: in the advanced economies, compared with today, 6 percent more of the working-age population uses it. But robots have not entered

yet, and tension stays between 2 and 4 points. Singapore, Israel, France, Ireland and Norway lead, because they are the countries that use the most AI.

In 2033 robots enter the advanced economies and the curve changes pace. In 2034 the United States, Germany, France, the United Kingdom and Israel are between 16 and 18 points. China, India and Brazil are at 3. Africa is close to zero.

In 2038 the advanced economies have lost almost 39 points of the work of the nineteen occupations to AI and robots, and median tension is 37. The emerging economies, where robots have just entered, are at 13. Low-income countries are at 2.

In 2046 the gap remains. The advanced economies range from 52 to 64 points, with a median of 56. The emerging ones from 30 to 44, with a median of 35. Low-income countries from 16 to 21, with a median of 18. The order never reverses. From 2030 onwards every advanced economy stands above every emerging economy, and every emerging economy above every low-income country: those with more AI and more robots have more tension.

Table 67 – Social tension by country, increase compared with January 2026. Work taken by AI and robots in the region, shared by AI use (Microsoft AI Economy Institute, Our World in Data) and amplified by political stability in 2024 (World Bank, Worldwide Governance Indicators). Data: aimodel, outputs/social_tension.json, the same as the Social Risk page.

Country	2030	2034	2038	2042	2046
Israel	3	18	42	58	64
United States	2	17	40	55	60
Greece	2	15	39	55	60
Germany	2	16	39	54	59
France	3	18	39	53	58
Italy	2	16	38	53	58
United Kingdom	3	17	38	52	57
Spain	3	17	38	53	57
South Korea	2	16	37	51	55
Japan	2	14	36	50	55
Singapore	4	18	35	48	52
United Arab Emirates	2	7	21	35	44
Russia	0	2	14	30	39
Egypt	0	3	14	30	38
Turkey	0	3	14	29	38
India	0	3	14	29	37
Indonesia	0	3	14	29	37
Brazil	1	3	14	28	36
Mexico	1	3	14	28	36
South Africa	1	3	14	28	36

Country	2030	2034	2038	2042	2046
China	0	3	13	27	35
Afghanistan	0	0	2	10	21
Yemen	0	1	3	10	20
Nigeria	0	0	2	9	20
Ethiopia	0	0	2	9	20

At the top in 2046 is Israel, with 64 points: it uses a lot of AI and is already tense today. Then come the United States and Greece with 60, Germany with 59, France, Belgium and Italy with 58. Japan is at 55: it uses less AI than the other advanced countries and is stable. Among the emerging economies the highest value is that of the United Arab Emirates, 44, because seven people of working age in ten already use AI. Russia, Egypt and Turkey are around 38–39, China at 35.

Africa is at the bottom. Not because it is calmer: Nigeria, Ethiopia, Yemen and Afghanistan start from high tension, and the amplifier works at its maximum. It is at the bottom because AI and robots arrive there last and find less exposed work. The sounding board amplifies what is there; if little work is replaced, the sound is weak too. This is why low-income countries stay close to one another, between 16 and 21 points, whatever their starting stability.

Within the same region the differences are small. The region decides, that is how much work is exposed and when robots enter. The country adjusts: with the AI it uses and with the fragility it carries.

The United States state by state

The United States is one country with fifty different labour markets. This is why the projection divides it state by state, using the employment of the nineteen occupations in each state from the Bureau of Labor Statistics, Occupational Employment and Wage Statistics, May 2024. The same formulas, with the weight of each state's occupations.

In 2038 the hardest hit is Utah, with 43 points, then Idaho with 42 and Washington, Florida and Georgia with 41. Utah has many call center agents and many software developers; Florida many retail clerks and call center agents, Georgia many retail clerks and warehouse workers. The least hit are Alaska, South Dakota and West Virginia, with 37, and the District of Columbia, with 35. In 2046, when robots weigh more, the states of factories and warehouses rise: Indiana, South Carolina, Georgia, Utah, Idaho and Tennessee reach 63. In Indiana more than one worker in four among the nineteen occupations works in a factory or a warehouse. New York stops at 57. The District of Columbia at 47: there almost one worker in three among the nineteen occupations is a lawyer, and there are hardly any factories.

California and Texas are at 40 in 2038 and at 61 in 2046. No state is left out: the gap between the hardest hit and the least hit, the District of Columbia aside, is no more than seven points.

China, Russia, Canada, Australia, Brazil and India are shown on the map by large city, because tension gathers where people live. For cities there are no data on occupations comparable across countries, so every city carries the value of its country.

What it is not

It is not a count of riots. Chapter 17 distinguished protest, unrest and violence, and that distinction stands. The tension number measures pressure; it does not say what form it will take.

The rules of the projection are these, written here only once. Starting tension amplifies by up to half as much again. Before robots enter, tension grows at a quarter of the speed. Robots enter a region when they take 3 percentage points of the work of the nineteen occupations. They are the author's rules, not estimated coefficients, and that is why they are declared. Changing them changes the levels. It does not change the order: first the advanced countries, then the emerging ones, last the low-income countries.

Some facts would refute the projection. If by 2034 robots had not taken at least 3 points of the work of the nineteen occupations in the advanced economies, the acceleration would have to be moved later. If AI use as measured by Our World in Data stopped growing, replacement would slow with it. If tension rose earlier and more in the countries that use less AI, the rule would be wrong. The data to check this come out every quarter and every year.

Where to see the projection

The projection is on the book's website, on the Social Risk page: https://thereplacemind.com/social.html. The map starts empty and fills month after month, from the period one chooses. The red dots are new AI users, the yellow ones the work taken by robots, the bubbles the tension of each country, each American state, each large city.

On the same site there are the simulator of the forecast by occupation, https://thereplacemind.com/simulator.html, and the economy model with its analyses, https://thereplacemind.com/

model.html. The code and the data are public: https://github.com/0xENX/aimodel_replication. The numbers in this chapter come from aimodel/analysis/social_tension.py, which repeats the page's calculations step by step.

Conclusions

Results and Open Questions

A machine that produces more and a growing number of people who cannot take part in what it produces. If there is one image that sums up this book, it is this.

In the model, world GDP grows compared with the world without AI: plus 1.3 percent in 2036, plus 3.2 in 2046. At the same time the share of income going to labor falls: minus 3 points in 2036, minus 5.9 in 2046. More efficiency does not on its own bring the benefit to those who live by their work.

The world's poorer half consumes 1 percent less in 2036 and 3.3 percent less in 2046. The richest ten percent, in 2036, 3.1 percent more. They say that growth of the total and the worsening of part of the population can go together. And they go together in nine futures out of ten.

The second problem is access. In the scenario in which retraining is slow, first entry of young people aged 15 to 24 in the advanced economies is 14.3 points lower in 2036. It is a precise risk: removing entry-level tasks makes it harder to gain experience, long before total employment shows the change. The American data of the last year point in the same direction.

The third is amplification. When those left without work weigh on wages and benefits, the loss of the world's poorer half in 2036 moves from 1.43 to 3.15 percent if retraining is slow, and from 2.64 to 7.24 if substitution persists. In the average scenario the difference is small. That is exactly why it is dangerous: it hides how much things can get worse when the transition jams.

Policies improve outcomes and leave choices open. With the same budget, the incentive to hire young people opens the entrance door more; reduced hours lower the total number of people without work more. The first risks paying for hires that would have happened anyway; the second asks those already working to give up a little income. And in all the experiments, protecting con-

sumption does not coincide with recovering work. No single ranking replaces the decision about what you want to achieve.

The concrete examples hold the numbers together. In the workshop, a faster draft becomes an advantage only after the fee, the checks and the capacity to deliver. In the library, a better service does not say who collects the fee or how the beginner learns. A bus compatible with the shift can make reachable a job that already exists. Between a technical possibility and a benefit a person can use there are many steps.

The assumptions are listed only once, at the end of chapter 6, the rules of the forecast by occupation in chapter 36 and those of social risk in chapter 37. The code is public so that anyone can redo the calculations and find the errors.

And above all: these are the results of the cautious scenario. Chapter 36 said openly what to expect occupation by occupation, robots included: by 2046 the office clerk loses almost six jobs in ten, the software developer four, the retail clerk almost seven, the factory worker six; where demand grows, as for nurses and electricians, jobs fall little. And it said which facts could refute it.

Chapter 37 carried the same forecast from work to social tension. More AI and robots, more people replaced, more tension, and first in the developed countries: in the advanced economies the curve accelerates from 2033, when robots enter, and by 2046 it reaches a median of 56 points, with Israel, the United States and Greece at the top. The emerging economies follow four years later and stop at 35, low-income countries at 18. Where a country is already tense the same replacement weighs more, but it never reverses the order.

The book's horizon ends in 2046. Its thesis does not. Artificial intelligence will reduce job opportunities in exposed activities, and reabsorption will remain incomplete if institutions do not change. Output growth will not on its own guarantee inclusion. Decisions

about who can enter work, how value is distributed and which services exist will change the outcome of this transition more than any technical innovation.

The risk is a richer economy in which a growing share of people struggle to earn an income and to count for something. It is for this risk that we must prepare. Not when the numbers have made it obvious to everyone, but now, while there is still time to build.

The simulator of the forecast by occupation, with the same data, the same rules and the same numbers as Table 65, is here: https://thereplacemind.com/simulator.html. The economy model with its analyses is here: https://thereplacemind.com/model.html. The projection of social risk, with the same numbers as Tables 66 and 67, is here: https://thereplacemind.com/social.html. The model code and data are here: https://github.com/0xENX/aimodel_replication.

Chapter Notes

For each of 4,000 combinations of assumptions, the simulations compare a world with the new wave of AI and one without it. Code, configurations and instructions for reproducing them are in the repository listed on the page "The Code Behind the Book".

Numerical examples described as teaching examples serve to explain a mechanism: they are neither observed data nor simulation results. The web addresses were consulted between October 10 and 16, 2026. Works cited in short form are given in full in the References.

Chapter 1. Output, Employment and Income Distribution

Tables 1 and 2 and the examples on the labor share use teaching numbers. Figure 1 is a diagram of the links between output, work and incomes; it does not assign them a strength.

Chapter 2. Tasks, Exposure and AI Adoption

Exposure measures: IMF, Gen-AI: Artificial Intelligence and the Future of Work (Cazzaniga and coauthors, 2024); ILO and NASK, Generative AI and Jobs: A Refined Global Index of Occupational Exposure (Gmyrek and co-authors, 2025). The percentages follow the sources' definitions and do not represent observed job losses. Table 5 and Figure 3 use a teaching example. The relationships for adoption, productive capacity and tasks are explained in Sections A1–A3 of the appendix.

Chapter 3. Participation, Unemployment and Discouragement

2025 data: Eurostat, labour force survey, series lfsa_pganws, Italy and EU-27, aggregation of five-year age groups (https://ec.europa.eu/eurostat/databrowser/view/lfsa_pganws/default/tabl

e); definitions: Eurostat metadata (https://ec.europa.eu/eurostat/cache/metadata/en/employ_esms.htm). Transitions between states: series lfsi_long_q, quarters 2024–2025, with no breakdown by age; a quarterly rate should not be multiplied by four. Figure 5: Stanford Digital Economy Lab, Canaries Dashboard, version of September 17, 2026, ADP sample, base November 2022 = 100; for the effect of AI every monthly change is weighted by its year (2026 and 2025 100%, 2024 50%, 2023 5%, 2022 1%) (https://digitaleconomy.stanford.edu/project/indicators/canaries-dashboard/). Tables 6–8 use teaching numbers.

Chapters 4–6. Labor Demand and Supply; Incomes and Consumption; the Model

The figures in Chapters 4–6 are calculated with the model. The world labor charts aggregate the three regions with the calibration's population weights: this is an approximation, not an official series of the labor force. Consumption by income group is used as an indicator of purchasing power. Figure 8 uses the age module in the slow-retraining scenario. The relationships for the labor market, households and demand are explained in Sections A4–A6 of the appendix.

Chapters 7–9. Flows, First Entry and Retraining

The age-based employment module is an extension developed for the book; its method is described in the repository (file FONTI_E_METODO.md). The figures use the module's simulations with the European Union profile as a partial reference for the advanced economies. Flows use each age group's potential supply in 2026 as their base; the first-entry rate uses eligible job seekers as its base. The "first entry" category is assigned by the model, not reconstructed from individual work histories. Hiring probabilities, barriers and retraining times are assumptions, not causal estimates. In Chapters 7–9 the link runs from the macroeconomic model to the module; Chapter

26 presents the experiment that makes it two-way. Chapter 8: Brynjolfsson, Li and Raymond, Generative AI at Work (NBER Working Paper 31161, 2023; Quarterly Journal of Economics, 2025). Chapter 9: the sensitivity analysis of the age module follows the method of Saltelli and coauthors (2010) and the Jansen estimator (1999).

Chapters 10–12. Value Added and Regions; Public Revenue and Debt; Calibration

The concepts of value added, national income, debt dynamics and Monte Carlo simulation are explained with teaching examples. References: OECD, Gross national income (https://www.oecd.org/en/data/indicators/gross-national-income.html); Adrian, Gaspar and Gourinchas, The Fiscal and Financial Risks of a High-Debt, Slow-Growth World, IMF blog, March 28, 2024; NIST Uncertainty Machine (https://uncertainty.nist.gov/). Widening new tasks up to 1.20 and the test with interdependent assumptions are additional experiments of the book.

Chapter 13. Public Policy Evaluation

Tables 28 and 29 and Figure 22: model simulations, 2,000 paired draws for the policy comparison.

Chapter 14. Aggregate Results and Heterogeneity

Table 30: model simulations. Table 31: model simulations and evidence cited in the text.

Chapter 15. Revising the Model

Tables 32 and 33 and Figure 23: comparison between versions of the model.

Chapter 16. Sensitivity Analysis

Method: Sobol' (2001); Saltelli and coauthors (2008).

Chapter 17. Social Cohesion and Institutions

Sources: Hadzi-Vaskov, Pienknagura and Ricci (2021); OECD, Survey on Drivers of Trust in Public Institutions, 2024 (https://doi.org/10.1787/9a20554b-en); NIST IR 8280 (Grother, Ngan and Hanaoka, 2019) and the NIST page on demographic differentials (https://pages.nist.gov/frvt/html/frvt_demographics.html); Regulation (EU) 2024/1689, Article 5. The example of flags is a teaching example.

Chapter 18. Model Limits and Empirical Testing

Sources: OECD, Job quality (https://www.oecd.org/en/topics/policy-issues/job-quality.html); Eurostat, Access to microdata (https://ec.europa.eu/eurostat/web/microdata/access); Istat, three-month longitudinal labour force files. Difference-in-differences is illustrated with an example.

Chapter 19. Review of the Evidence

Figure 26: coefficients from Acemoglu and Restrepo (2019). Additional checks: Acemoglu (2025); Brynjolfsson, Li and Raymond (2025); Bick, Blandin and Deming, The Rapid Adoption of Generative AI (NBER Working Paper 32966); Brynjolfsson, Chandar and Chen (2026). The review is not a systematic review.

Chapters 20–22. Tax Bases, Scenarios and Interventions

Tax rates and public revenue configurations are assumptions for unidentified economies, not rules in force.

Chapter 23. The Causal Structure of the Model

The model's two summary relationships are reported in Section A10 of the appendix.

Chapter 24. Training and Learning

Bastani and coauthors, Generative AI without guardrails can harm learning, PNAS, 2025 (https://doi.org/10.1073/pnas.2422633122). De Simone and coauthors, From Chalkboards to Chatbots, World Bank Policy Research Working Paper 11125, 2025 (https://doi.org/10.1596/1813-9450-11125). The comparison of the three policies with the same budget (0.7% of GDP a year) is an additional experiment of the book; its costs are not carried into the macroeconomic accounts.

Chapter 25. Job Quality and Pay

Cazes, Hijzen and Saint-Martin, Measuring and Assessing Job Quality, OECD, 2015 (https://doi.org/10.1787/5jrp02kjw1mr-en). Lane, Williams and Broecke, The impact of AI on the workplace, OECD, 2023 (https://doi.org/10.1787/ea0a0fe1-en). Rani, Pesole and González Vázquez, Algorithmic management practices in regular workplaces, JRC and ILO, 2024.

Chapter 26. AI Adoption in Small Firms

Eurostat, Use of artificial intelligence in enterprises, 2025 data. OECD, Generative AI and the SME Workforce: New Survey Evidence, 2025 (https://doi.org/10.1787/2d08b99d-en). The values in the figures are taken from the publications. The experiment with the two-way link between employment, wages and benefits is an extension of the book: both worlds, with and without AI, are calculated with the same structure.

Chapter 27. Capital Ownership and Concentration

Autor, Dorn, Katz, Patterson and Van Reenen, The Fall of the Labor Share and the Rise of Superstar Firms, Quarterly Journal of Economics, 2020.

Chapter 28. AI and Health Services

Lång and coauthors, MASAI, Lancet Oncology, 2023 (https://doi.org/10.1016/S1470-2045(23)00298-X). Gommers and coauthors, Lancet, 2026 (https://doi.org/10.1016/S0140-6736(25)02464-X). Goh and coauthors, JAMA Network Open, 2024 (https://doi.org/10.1001/jamanetworkopen.2024.40969).

Chapter 29. The Energy Use of AI

IEA, Key Questions on Energy and AI, April 16, 2026. Historical estimates and the central projection are kept distinct.

Chapter 30. Mobility and Territorial Differences

Acemoglu and Restrepo (2020); Dauth, Findeisen, Suedekum and Woessner (2021); Chetty, Hendren and Katz, The Effects of Exposure to Better Neighborhoods on Children, American Economic Review, 2016 (https://doi.org/10.1257/aer.20150572). Changes 2022–2025 for Italy and the EU: descriptive processing of the Eurostat series lfsa_pganws; yearly pace weighted by year as in chapter 3.

Chapter 31. Impact Evaluation of Interventions

Katz, Roth, Hendra and Schaberg, Why Do Sectoral Employment Programs Work?, Journal of Labor Economics, 2022. The randomized experiment, non-response and cost per additional person employed are teaching examples.

Chapter 32. Results and the Limits of Knowledge

The outline of empirical questions is a proposal for observation.

Chapters 33–35. Summary

These chapters bring together results already presented in Chapters 8, 13, 14, 22, 24 and 26. The cases of the workshop, the library, the café and transport are hypothetical examples.

Chapter 36. The Last Forecast

The chapter is the book's forecast, based on the earlier chapters and on the following sources. Table 65 comes from aimodel (aimodel/analysis/occupations.py, outputs/occupations.json), with the rules listed in the chapter and sources [36.6]–[36.8]; the simulator uses the same data. The intervals 2026–2030, 2030–2036 and 2036–2046 organize the forecast.

[36.1] Brynjolfsson, E., Chandar, B. and Chen, R. (2026). Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence. Stanford Digital Economy Lab, revision of August 12, 2026.

[36.2] Massenkoff, M. and McCrory, P. (2026). Labor market impacts of AI: A new measure and early evidence. Anthropic, March 5, 2026. Research produced by a company in the industry.

[36.3] Gmyrek, P., Berg, J. and coauthors (2025). Generative AI and jobs: A refined global index of occupational exposure. ILO Working Paper 140.

[36.4] METR (2026). Time Horizon 1.1. January 29, 2026.

[36.5] Bustos, P., Castro-Vincenzi, J. M., Monras, J. and Ponticelli, J. (2019, revised 2022). Industrialization without innovation. NBER Working Paper 25871. Studies reallocation between sectors linked to agricultural productivity in Brazil.

[36.6] Eloundou, T., Manning, S., Mishkin, P. and Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. Science, 384, 1306–1308. Gamma measure rated by human experts.

[36.7] Frey, C. B. and Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? Technological Forecasting and Social Change, 114, 254–280 (working paper 2013).

[36.8] U.S. Bureau of Labor Statistics (2026). Employment Projections 2025–35, table Occupational projections and characteristics, September 21, 2026.

Chapter 37. Social Risk

The chapter is a projection with rules declared in the chapter itself. The numbers in Tables 66 and 67 come from aimodel (outputs/social_tension.json, written by aimodel/analysis/social_tension.py), the same as the Social Risk page of the website; the work taken by AI and robots comes from the nineteen occupations of Table 65 and from the sources of chapter 36.

[37.1] World Bank (2025). Worldwide Governance Indicators, Political Stability and Absence of Violence/Terrorism, 2024 estimate (GOV_WGI_PV.EST). CC BY 4.0 licence.

[37.2] Microsoft AI Economy Institute, via Our World in Data (2026). Estimated share of people using generative AI: share of the working-age population, quarterly from June 2025 to March 2026. CC BY 4.0 licence.

[37.3] U.S. Bureau of Labor Statistics (2025). Occupational Employment and Wage Statistics, May 2024, state estimates.

[37.4] International Monetary Fund. World Economic Outlook, country groups: advanced economies. World Bank, country classification by income level.

Appendix

Sections A1–A10: relationships of the model. Sections A11 and A12 concern extensions and teaching calculations of the book.

Cover

Illustration generated with artificial intelligence tools. The use of artificial intelligence in preparing the text and the code is disclosed in the Introduction.

References

List of the works cited in the text and used to build the model, with titles in their original language. Data sources and studies referred to for a single passage are also indicated in the chapter notes, with their web addresses.

Acemoglu, D. (2025). The simple macroeconomics of AI. Economic Policy, 40(121), 13–58. NBER Working Paper 32487.

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

Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. Journal of Economic Perspectives, 33(2), 3–30. doi:10.1257/jep.33.2.3

Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. Journal of Political Economy, 128(6), 2188–2244. doi:10.1086/705716

Acemoglu, D. and Restrepo, P. (2022). Tasks, automation, and the rise in US wage inequality. Econometrica, 90(5), 1973–2016. doi:10.3982/ECTA19815

Allen, R. C. (2009). Engels' pause: Technical change, capital accumulation, and inequality in the British industrial revolution. Explorations in Economic History, 46(4), 418–435.

Autor, D., Chin, C., Salomons, A. and Seegmiller, B. (2024). New frontiers: The origins and content of new work, 1940–2018. Quarterly Journal of Economics, 139(3), 1399–1465. doi:10.1093/qje/qjae008

Autor, D. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US labor market. American Economic Review, 103(5), 1553–1597. doi:10.1257/aer.103.5.1553

Banerjee, A., Faye, M., Krueger, A., Niehaus, P. and Suri, T. (2023). Universal basic income: Short-term results from a long-term experiment in Kenya. NBER conference paper.

Bank of Canada (2019). Exploring wage Phillips curves in advanced economies. Staff Discussion Paper 2019-8.

Brollo, F., Dabla-Norris, E., de Mooij, R., Garcia-Macia, D., Hanappi, T., Liu, L. and Nguyen, A. D. M. (2024). Broadening the gains from generative AI: The role of fiscal policies. IMF Staff Discussion Note SDN/2024/002.

Brynjolfsson, E., Chandar, B. and Chen, R. (2026). Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence. Stanford Digital Economy Lab, revision of August 12, 2026.

Brynjolfsson, E., Li, D. and Raymond, L. (2025). Generative AI at work. Quarterly Journal of Economics. Earlier version: NBER Working Paper 31161, 2023.

Brynjolfsson, E., Rock, D. and Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. American Economic Journal: Macroeconomics, 13(1), 333–372. doi:10.1257/mac.20180386

Bustos, P., Castro-Vincenzi, J. M., Monras, J. and Ponticelli, J. (2019, revised 2022). Industrialization without innovation. NBER Working Paper 25871. doi:10.3386/w25871

Carroll, C., Slacalek, J., Tokuoka, K. and White, M. N. (2017). The distribution of wealth and the marginal propensity to consume. Quantitative Economics, 8(3), 977–1020.

Cazzaniga, M., Jaumotte, F., Li, L., Melina, G., Panton, A. J., Pizzinelli, C., Rockall, E. and Tavares, M. M. (2024). Gen-AI: Artificial intelligence and the future of work. IMF Staff Discussion Note SDN/2024/001.

Comin, D. and Hobijn, B. (2010). An exploration of technology diffusion. American Economic Review, 100(5), 2031–2059. doi:10.1257/aer.100.5.2031

European Commission (2025). Guidelines on prohibited artificial intelligence practices established by Regulation (EU) 2024/1689 (AI Act). February 2025.

Dauth, W., Findeisen, S., Suedekum, J. and Woessner, N. (2021). The adjustment of labor markets to robots. Journal of the European Economic Association, 19(6), 3104–3153. doi:10.1093/jeea/jvab012

David, P. A. (1990). The dynamo and the computer: An historical perspective on the modern productivity paradox. American Economic Review, 80(2), 355–361.

Eloundou, T., Manning, S., Mishkin, P. and Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. Science, 384(6702), 1306–1308. doi:10.1126/science.adj0998

Federal Reserve Bank of San Francisco (2017). Has the wage Phillips curve gone dormant? Economic Letter 2017-30.

Fisher, J. D., Johnson, D. S., Smeeding, T. M. and Thompson, J. P. (2020). Estimating the marginal propensity to consume using the distributions of income, consumption, and wealth. Journal of Macroeconomics, 65.

Giupponi, G. and Landais, C. (2023). Subsidizing labour hoarding in recessions: The employment and welfare effects of short-time work. Review of Economic Studies, 90(4), 1963–2005.

Gmyrek, P., Berg, J. and coauthors (2025). Generative AI and jobs: A refined global index of occupational exposure. ILO Working Paper 140. doi:10.54394/HETP0387

Godley, W. and Lavoie, M. (2007). Monetary Economics: An Integrated Approach to Credit, Money, Income, Production and Wealth. Palgrave Macmillan. doi:10.1057/9780230626546

Grother, P., Ngan, M. and Hanaoka, K. (2019). Face Recognition Vendor Test (FRVT) Part 3: Demographic effects. NIST IR 8280.

Hadzi-Vaskov, M., Pienknagura, S. and Ricci, L. A. (2021). The macroeconomic impact of social unrest. IMF Working Paper 2021/135.

Hamermesh, D. S. (1993). Labor Demand. Princeton University Press.

International Energy Agency (2025). Energy and AI. Paris: IEA.

International Energy Agency (2026). Key Questions on Energy and AI. April 16, 2026.

International Federation of Robotics (2025). World Robotics 2025: Industrial Robots.

International Labour Organization (2024, 2025). World Employment and Social Outlook: September 2024 update; May 2025 update.

International Monetary Fund (2025, 2026). Fiscal Monitor, October 2025 and April 2026.

International Monetary Fund (2026). Annual Report 2026. AI: Deployment and Disruption.

Jansen, M. J. W. (1999). Analysis of variance designs for model output. Computer Physics Communications, 117(1–2), 35–43.

Jones, D. and Marinescu, I. (2022). The labor market impacts of universal and permanent cash transfers: Evidence from the Alaska Permanent Fund. American Economic Journal: Economic Policy, 14(2), 315–340.

Karabarbounis, L. and Neiman, B. (2014). The global decline of the labor share. Quarterly Journal of Economics, 129(1), 61–103. doi:10.1093/qje/qjt032

Katz, L. F., Roth, J., Hendra, R. and Schaberg, K. (2022). Why do sectoral employment programs work? Lessons from WorkAdvance. Journal of Labor Economics, 40(S1), S249–S291.

Kela and Ministry of Social Affairs and Health, Finland (2020). Results of the basic income experiment.

Korinek, A. and Suh, D. (2024). Scenarios for the transition to AGI. NBER Working Paper 32255.

Lichter, A., Peichl, A. and Siegloch, S. (2015). The own-wage elasticity of labor demand: A meta-regression analysis. European Economic Review, 80, 94–119.

Massenkoff, M. and McCrory, P. (2026). Labor market impacts of AI: A new measure and early evidence. Anthropic, March 5, 2026.

McKinsey Global Institute (2023). The economic potential of generative AI: The next productivity frontier. June 2023.

METR (2026). Time Horizon 1.1. January 29, 2026.

Oberfield, E. and Raval, D. (2021). Micro data and macro technology. Econometrica, 89(2), 703–732. doi:10.3982/ECTA12807

OECD (2025). Pensions at a Glance 2025. Paris: OECD Publishing.

OECD (2025). Revenue Statistics 2025. Paris: OECD Publishing.

Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M. and Tarantola, S. (2008). Global Sensitivity Analysis: The Primer. Wiley.

Saltelli, A., Annoni, P., Azzini, I., Campolongo, F., Ratto, M. and Tarantola, S. (2010). Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index. Computer Physics Communications, 181(2), 259–270.

Sobol', I. M. (2001). Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. Mathematics and Computers in Simulation, 55(1–3), 271–280.

European Union (2024). Regulation (EU) 2024/1689 on artificial intelligence (AI Act).

United Nations, Department of Economic and Social Affairs (2024). World Population Prospects 2024.

World Economic Forum (2025). Future of Jobs Report 2025.

World Inequality Lab (2022, 2026). World Inequality Report 2022; World Inequality Report 2026.

Writers Guild of America (2023). Minimum Basic Agreement: provisions on artificial intelligence.

Other sources

Facts and historical cases recalled in the narrative of the chapters, which are not used to build the model. They are listed in the order in which they appear, with the source against which they can be checked.

Chapter 3. Statute of Artificers (1563), 5 Eliz. I c. 4, England. Wallis, P. (2008). Apprenticeship and training in premodern England. Journal of Economic History, 68(3), 832–861.

Chapter 4. Royal Swedish Academy of Sciences (2010). Markets with search frictions. Scientific background on the Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2010 (Diamond, Mortensen, Pissarides). U.S. Bureau of Labor Statistics (2022). Number of unemployed people per job opening is 0.5 in March 2022. The Economics Daily, https://www.bls.gov/opub/ted/2022/number-of-unemployed-people-per-job-opening-is-0-5-in-march-2022.htm.

Chapter 5. Meyer, S. (1981). The Five Dollar Day: Labor Management and Social Control in the Ford Motor Company, 1908–1921. Albany: State University of New York Press. Robinson, J. (1966). Kalecki and Keynes. In Problems of Economic Dynamics and Planning: Essays in Honour of Michal Kalecki. Oxford: Pergamon Press.

Chapter 6. Lind, J. (1753). A Treatise of the Scurvy. Edinburgh: Sands, Murray and Cochran.

Chapter 7. Jacobson, L. S., LaLonde, R. J. and Sullivan, D. G. (1993). Earnings losses of displaced workers. American Economic Review, 83(4), 685–709.

Chapter 8. Kahn, L. B. (2010). The long-term labor market consequences of graduating from college in a bad economy. Labour Economics, 17(2), 303–316. Oreopoulos, P., von Wachter, T. and Heisz, A. (2012). The short- and long-term career effects of graduating in a recession. American Economic Journal: Applied Economics, 4(1), 1–29.

Chapter 9. Schochet, P. Z., D'Amico, R., Berk, J., Dolfin, S. and Wozny, N. (2012). Estimated Impacts for Participants in the Trade Adjustment Assistance (TAA) Program under the 2002 Amendments. Mathematica Policy Research and Social Policy Research Associates for the U.S. Department of Labor.

Chapter 10. Central Statistics Office (2016). National Income and Expenditure 2015, 12 July 2016, Cork. Economic Statistics Review Group (2016). Report of the Economic Statistics Review Group, December 2016, Dublin (proposal of modified gross national income, GNI*). The expression "leprechaun economics" was used by Paul Krugman in a public comment on 12 July 2016.

Chapter 11. Delaney, K. J. (2017). The robot that takes your job should pay taxes, says Bill Gates. Quartz, 17 February 2017, https://qz.com/911968/bill-gates-the-robot-that-takes-your-job-should-pay-taxes/. European Parliament (2017). Resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics, 2015/2103(INL).

Chapter 12. Reinhart, C. M. and Rogoff, K. S. (2010). Growth in a time of debt. American Economic Review, 100(2), 573–578. Herndon, T., Ash, M. and Pollin, R. (2014). Does high public debt consistently stifle economic growth? A critique of Reinhart and Rogoff. Cambridge Journal of Economics, 38(2), 257–279. Metropolis, N. and Ulam, S. (1949). The Monte Carlo method. Journal of the American Statistical Association, 44(247), 335–341.

Metropolis, N. (1987). The beginning of the Monte Carlo method. Los Alamos Science, 15, 125–130. Engels, F. (1845). Die Lage der arbeitenden Klasse in England. Leipzig: Otto Wigand (English translation: The Condition of the Working Class in England).

Chapter 13. Alaska Department of Revenue, Permanent Fund Dividend Division. Historical dividend amounts, since 1982. Kela (2020). The Basic Income Experiment 2017–2018 in Finland: final results (2,000 unemployed people selected by lottery, 560 euros a month).

Chapter 15. Feynman, R. P. (1974). Cargo cult science. Engineering and Science, 37(7), 10–13 (commencement address at the California Institute of Technology).

Chapter 16. Sobol', I. M. (1993). Sensitivity estimates for nonlinear mathematical models. Mathematical Modelling and Computational Experiments, 1(4), 407–414.

Chapter 17. Parlementaire ondervragingscommissie Kinderopvangtoeslag (2020). Ongekend onrecht. Tweede Kamer der Staten-Generaal, 17 December 2020. Autoriteit Persoonsgegevens (2020). Belastingdienst/Toeslagen: de verwerking van de nationaliteit van aanvragers van kinderopvangtoeslag, July 2020.

Chapter 18. Card, D. and Krueger, A. B. (1994). Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania. American Economic Review, 84(4), 772–793. Royal Swedish Academy of Sciences (2021). Press release: The Prize in Economic Sciences 2021 (Card; Angrist and Imbens).

Chapter 19. Solow, R. M. (1987). We'd better watch out. New York Times Book Review, 12 July 1987, p. 36.

Chapter 20. OECD/G20 Inclusive Framework on BEPS (2021). Statement on a Two-Pillar Solution to Address the Tax Challenges Arising from the Digitalisation of the Economy, 8 October 2021.

Chapter 21. Wack, P. (1985). Scenarios: Uncharted waters ahead. Harvard Business Review, 63(5), 72–89.

Chapter 23. Phillips, A. W. (1950). Mechanical models in economic dynamics. Economica, 17(67), 283–305.

Chapter 24. Plato, Phaedrus, 274c–275b.

Chapter 25. Court of Bologna, labor section, order of 31 December 2020 (Filcams, Nidil and Filt CGIL Bologna v. Deliveroo Italy). Commentary: Fassina, L. (2021). Lavoro Diritti Europa, 1/2021.

Chapter 26. Istat (2023). Annuario statistico italiano 2023, chapter 14, Imprese (microenterprises with fewer than 10 employees accounting for 95.2 percent of active firms in 2021).

Chapter 27. Piketty, T. (2013). Le Capital au XXIe siècle. Paris: Seuil (English translation: Capital in the Twenty-First Century, Cambridge, MA: Harvard University Press, 2014).

Chapter 28. Leveson, N. G. and Turner, C. S. (1993). An investigation of the Therac-25 accidents. IEEE Computer, 26(7), 18–41.

Chapter 29. Jevons, W. S. (1865). The Coal Question. London: Macmillan.

Chapter 30. Kain, J. F. (1968). Housing segregation, Negro employment, and metropolitan decentralization. Quarterly Journal of Economics, 82(2), 175–197.

Chapter 31. Bloom, H. S., Orr, L. L., Bell, S. H., Cave, G., Doolittle, F., Lin, W. and Bos, J. M. (1997). The benefits and costs of JTPA Title II-A programs: Key findings from the National Job Training Partnership Act Study. Journal of Human Resources, 32(3), 549–576.

Chapter 34. U.S. Bureau of Labor Statistics (2015). Long-term unemployment in the United States. Spotlight on Statistics (record 45.1 percent of the unemployed out of work for 27 weeks or more in the second quarter of 2010).

Appendix: The Formulas Explained Step by Step

Sections

A1 The Diffusion of Technology 373
A2 The Capacity to Produce 374
A3 Human Tasks and Distribution 376
A4 Wages and Labor Demand 378
A5 From Household Income to Wealth 380
A6 From Demand to Actual Output 382
A7 The Assumptions That Matter in Sensitivity 384
A8 Public Revenue and Debt 385
A9 Checking the Counterparts 386
A10 The Arrows of the Task Model 388
A11 Sharing Working Hours 391
A12 Concentration and the Aggregate Labor Share 392

This appendix makes it possible to understand what the formulas do even without carrying out the calculations. Each section starts from a question, keeps the original relationships and guides the reading. The examples added serve to illustrate a mechanism: they are neither observed data nor new simulations. The explanations do not change the parameters or the results of the book.

An Independent Reading Path

The appendix can be read from the beginning or by opening the section that answers one's question. Each section reports the related chapters, the necessary definitions, the original relationships, an example and the limits. Cross-references to other sections allow deeper reading while keeping the link with the chapters accessible. The examples are illustrative and add no results to the simulations.

The Map of Links

A1 describes the adoption of the technology. A3 follows human tasks and A2 the capacity to produce. A4 links tasks, output and wages to labor demand. A5 follows incomes, consumption and wealth; A6 links spending to actual output. A8 reconstructs the public account and A9 checks the counterparts. The relationships influence one another: this map is a guide to the calculation, not causal proof.

A7 makes it possible to read the sensitivity of the results to the assumptions. A10 keeps the model's summary formulation, to be distinguished from the revised relationships. A11 looks more closely at sharing working hours; A12 distinguishes the concentration of sales from the distribution of value added.

How to Read the Symbols

The subscript t denotes the current year and t−1 the previous one; a subscript g distinguishes an income group. A product means that the factors act together. min picks the smaller value, max the larger. Σ adds up the components indicated; Δ measures a change. One minus a share expresses the part that remains. The same symbol can change meaning between modules: E denotes exposure in the macroeconomic relationships and the employed in those by age. Local definitions therefore take precedence.

A1 The Diffusion of Technology

Chapters 2 and 6.

No preliminary reading. Exposure means technical possibility, adoption means actual use.

How quickly does a technical capability become actual use?

The Mechanism and the Original Relationships

AI adoption in exposed, substitutable tasks follows a logistic curve, normalized to zero in 2026. θ denotes its upper limit, m the central point and k the slope. The central point arrives later in the emerging and low-income regions, according to the regional lags. Exposure grows linearly up to a multiple x of the initial value in 2046 and does not exceed 0.95.

$$a_t = \theta \frac{F(t) - F(2026)}{1 - F(2026)}$$

$$F(t) = \frac{1}{1 + e^{-k(t-m)}}$$

$$E_t = min(0.95;\ E_0 [1 + (x - 1)\frac{t - 2026}{20}])$$

Normalizing the curve means subtracting the level it would already have in 2026. The value a_t therefore represents the additional

modeled wave of adoption, not a measure of all the automation already present in the economy. The ceiling on exposure is a constraint of the model.

Reading the Formula One Step at a Time

Read F as the progress along an S-shaped curve. Subtracting the 2026 level removes what was already present at the starting point; the division rescales and θ sets the limit. In the exposure relationship, min picks the smaller of the calculated value and the ceiling of 0.95. A higher limit widens possible adoption; a higher k concentrates diffusion in fewer years.

An Example to Get Your Bearings

Imagine one hundred exposed activities and an adoption limit of sixty percent. The model allows sixty to be adopted once diffusion is complete; not that sixty workers necessarily lose their jobs. If the central point moves from 2035 to 2040, the path shifts forward. These are illustrative numbers, not results of the book.

The Limit to Remember

The curve is not an observed timetable and does not turn exposure into layoffs.

A2 The Capacity to Produce

Chapters 2 and 23.

Γ is the human labor content; δ is the rate of substitution and ϱ the rate of task recovery. E denotes exposure, a adoption and φ the substitutable part. These factors are examined in A1 and A3, but here it is enough to distinguish their roles.

Which contributions does the growth of potential output depend on?

The Mechanism and the Original Relationships

Baseline growth $\bar{g}$ combines with the productivity gain from AI. π is the cost saving on automated tasks; μ is the gain in tasks where AI assists human labor. The original formula is shown below, spread over three lines belonging to the same relationship.

$$Y^s_t = Y^s_{t-1}\,[1 + \bar{g} + \Gamma_{t-1}\,\delta_t\,\pi$$

$$+ \mu(1 - \varphi)\,max(0;\, E_t\,a_t - E_{t-1}\,a_{t-1})\,\Gamma_{t-1}$$

$$+ \frac{1}{2}\,\pi\,\Gamma_{t-1}\,\rho_t\,]$$

Y^s denotes potential output. Actual output Y can be lower: the demand closure is explained in A6. The creation of new tasks also contributes to output in the relationship implemented; ψ is therefore not just a parameter of individual re-employment.

Reading the Formula One Step at a Time

The bracket is a growth factor applied to the previous year's capacity. The first AI contribution combines the labor involved, substitution and the cost saving. The second combines the increase in adopted exposure with the help given to non-substitutable tasks. The third weights new tasks. max with zero prevents a negative increase from being counted as a new positive contribution.

An Example to Get Your Bearings

If initial capacity is one hundred and, in an example, baseline growth is two points and the three AI contributions together are one point, the new capacity is one hundred and three. This does not mean that all one hundred and three units are sold.

The Limit to Remember

This relationship describes productive capacity, not the income of every household or employment.

A3 Human Tasks and Distribution

Chapters 2 and 23.

E is exposure and a adoption; their product describes the exposed part actually adopted. s^L is the labor income share, distinct from the coefficient s for recovery through bargaining. A1 looks more closely at diffusion.

How do substitution, new tasks and the labor income share combine?

The Mechanism and the Original Relationships

In the model's equations, the content of human tasks in production falls through substitution and rises through the creation of new tasks. To introduce the structure we use the following simplified version:

$$\Gamma_t = \Gamma_{t-1} \times (1 - \delta_t + \rho_t)$$

Γ represents the task content attributed to labor in the model. δ represents the substitution rate for the period; ϱ the rate of creation or recovery. The full version also includes the evolution not attributed to AI and defines these rates as functions of exposure, adoption and lags. The formula is not a direct count of people laid off and hired.

An important assumption of the model links the content of human tasks to the labor income share through a particular production structure. It is a theoretical choice, not an identity valid for every economy. Assessing it requires reading the relationships for tasks, production and the labor market together. The equations

must state how a technical change becomes an economic change; changing the production structure would require a separate experiment, not just changing a number.

The simplified relationship isolates substitution and recovery. The full relationships add the baseline erosion not attributed to AI and the comparison with a path without the new technological wave.

Substitution and the creation of new tasks change the labor content of production. d_0 is the erosion not due to AI; φ is the substitutable fraction; ψ denotes new tasks per replaced task; ℓ is the lag with which they appear. δ_t measures substitution and ϱ_t recovery. Workers can recover through bargaining a fraction s of the fall in labor content caused by AI.

$$\delta_t = \varphi \, max(0;\, E_t \, a_t - E_{t-1} \, a_{t-1})$$

$$\rho_t = \psi \, \varphi \, max(0;\, E_{t-\ell} \, a_{t-\ell} - E_{t-1-\ell} \, a_{t-1-\ell})$$

$$\Gamma_t = \Gamma_{t-1} \, (1 - d_0 - \delta_t + \rho_t)$$

$$\Gamma^n{}_t = \Gamma^n{}_{t-1} \, (1 - d_0)$$

$$s^L{}_t = \Gamma_t + s \, max(0;\, \Gamma^n{}_t - \Gamma_t)$$

Γ measures the labor content of tasks and Γ^n the no-AI path. The summary relationship for s^L describes the labor income share in the model's summary specification. Section A4 reports the revised relationship with relative wage and demand elasticity. The two formulations have different functions and are not independent observed measures.

Reading the Formula One Step at a Time

In the simplified version, one starts from the previous value, subtracts substitution and adds recovery. The full version also sub-

tracts the erosion not attributed to AI. Recovery looks at adoption in earlier years: ℓ is the waiting time. The summary relationship for the labor share adds the part recovered through bargaining. The more detailed wage relationship is in A4.

An Example to Get Your Bearings

With initial labor content equal to one, baseline erosion of 1 percent, substitution of 5 percent and recovery of 2 percent, the next value is 0.96. These are teaching values, outside the book's calibration: they indicate a 4 percent fall in labor content, not four layoffs out of a hundred.

The Limit to Remember

Tasks, people and incomes are different quantities. The relationship between them is an economic assumption, not a simple arithmetic conversion.

A4 Wages and Labor Demand

Chapters 4 and 15.

Y is output, Γ the human labor content and d_0 the baseline annual erosion. s^L is the labor income share, s the bargaining coefficient and Γ^n the no-AI reference. g^{pot} is the growth of potential output; u is unemployment. A2 and A3 look more closely at output and tasks.

Why does a wage cut not automatically bring back all the jobs?

The Mechanism and the Original Relationships

Labor demand responds to the wage cost, given output and tasks. The elasticity ε measures the size of this response: when it is below one, a proportional wage cut recovers only part of the demand lost.

L^s is labor supply; u* structural unemployment; β the pass-through to wages of productivity growth above baseline; $\varkappa$ the slope of the wage Phillips curve. g^y is the trend growth of output per worker without AI and g^w the wage trend, net of erosion d_0. $\bar{w}$ is the wage in the paired no-AI simulation. ε is the elasticity of labor demand with respect to the wage, at constant output.

$$w_t = w_{t-1}\,[1 + g^{w}{}_t + \beta(g^{pot}{}_t - g^{y}{}_t)$$

$$- \kappa(u_{t-1} - u^*)\,]$$

$$1 + g^{w}{}_t = (1 + g^{y}{}_t)(1 - d_0)$$

$$N^{d}{}_t = \frac{\Gamma_t Y_t}{\bar{w}_t}\left(\frac{w_t}{\bar{w}_t}\right)^{-\varepsilon}$$

$$N_t = min(N^{d}{}_t;\ L^{s}{}_t(1 - u^*))$$

$$s^{L}{}_t = \Gamma_t\left(\frac{w_t}{\bar{w}_t}\right)^{1-\varepsilon} + s\,max(0;\ \Gamma^{n}{}_t - \Gamma_t)$$

$$u_t = 1 - \frac{N_t}{L^{s}{}_t}$$

N^d is the number of workers demanded by firms at the current wage; N is actual employment. If demand exceeds supply, the wage rises to clear the market. The condition is:

$$w = \frac{s^{L} Y}{L^{s}(1 - u^*)}$$

In the no-AI simulation w coincides with $\bar{w}$ and labor demand is ΓY/w, as in the first edition. In the revised edition, a fall in the wage restores only part of the lost demand when ε is below one. The bargaining component changes the share of income attributed to labor; in the relationship for N^d it is not added as automatic job creation.

Reading the Formula One Step at a Time

The ratio between the current wage and the no-AI wage measures relative cost. The negative exponent makes demand grow when this ratio falls. min limits employment to what firms and supply together allow. The unemployment formula divides employment by supply and subtracts the ratio from one. The wage-clearing relationship applies to the condition of demand exceeding supply.

An Example to Get Your Bearings

With an elasticity of 0.3, a 1 percent wage cut increases labor demand by about 0.3 percent for a small change, holding output and other conditions fixed. The cost benefit is therefore not a complete recovery of employment.

The Limit to Remember

The wage response depends on the assumed elasticity; the effects of wages on consumption also require A5 and A6.

A5 From Household Income to Wealth

Chapters 5 and 10.

Y is output, s^L the labor income share and W wealth. Group g receives part of labor income and part of capital income. The tax rates τL and τK concern labor and capital respectively; h in the final sum denotes all groups, not the hours of Section A11.

How can today's distribution change tomorrow's incomes?

The Mechanism and the Original Relationships

The aggregate labor share is not enough to determine the income of each household. Ownership of capital and transfers contribute to distribution across groups.

λg is the share of labor income attributed to group g, ωg its share of wealth, τ the tax rates, ι the share of profits retained by firms, Tr transfers and cg the propensity to consume. Transfers include pensions, benefits, interest and policy measures. The group shares are regional.

$$Y^d_{g,t} = \lambda_g (1 - \tau_L) s^{\mathrm{L}}{}_{\mathrm{t}} Y_{\mathrm{t}}$$

$$+ \omega_{g,t} (1 - \tau_K)(1 - \iota)(1 - s^{\mathrm{L}}{}_{\mathrm{t}}) Y_{\mathrm{t}} + Tr_{g,t}$$

$$C_{g,t} = c_g Y^d_{g,t}$$

$$W_{g,t+1} = W_{g,t} + (Y^d_{g,t} - C_{g,t})$$

$$+ \omega_{g,t} (1 - \tau_K) \iota (1 - s^{\mathrm{L}}{}_{\mathrm{t}}) Y_{\mathrm{t}}$$

$$\omega_{g,t} = \frac{W_{g,t}}{\Sigma_h W_{h,t}}$$

$Y^{d}g$ is the group's disposable income and Wg its wealth. Distributed capital income enters disposable income; retained profits contribute to the accumulation of wealth. This distinction allows distribution to evolve over time. The model does not follow differences between people within each group.

Reading the Formula One Step at a Time

The first relationship adds up labor income, distributed capital income and transfers. The second applies the propensity to consume. The third updates wealth with saving and retained profits. The last divides the group's wealth by total wealth: the ownership share can thus change over time. Tax rates reduce the components to which they apply.

An Example to Get Your Bearings

A group with disposable income of one hundred and a propensity to consume of 0.8 consumes eighty and saves twenty. If five of retained profits are also attributed to it, its wealth grows by twenty-five in the simplified account. The example does not estimate returns or real households.

The Limit to Remember

A group average does not represent every individual. The model does not reconstruct all internal differences.

A6 From Demand to Actual Output

Chapters 5 and 6.

Y^s is productive capacity, Y^d demand and Y actual output. s^L is the labor income share. c^y is consumption per unit of output; c^{tr} is consumption linked to transfers. The noAI reference is the paired simulation without the new wave of AI. A2 and A5 look more closely at capacity and consumption.

Who buys what the economy is able to produce?

The Mechanism and the Original Relationships

Demand links household spending to actual output. Investment and government purchases take part in the same account, together with the transfers that feed consumption.

Investment responds to the profit share with elasticity η; i is the initial investment share; G is government purchases and τC the consumption tax. Effective demand solves a linear fixed point. Output responds to any demand shortfall with coefficient ν, relative to the ratio in the no-AI simulation, in which the economy is assumed to operate at potential.

$$I_t = i\, Y^s{}_t \left(\frac{1 - s^L{}_t}{1 - s^L{}_{2026}}\right)^{\eta}$$

$$Y^d{}_t = \frac{(1 - \tau_C)c^{tr}{}_t + I_t + G_t}{1 - (1 - \tau_C)c^y{}_t}$$

$$Y_t = Y^s{}_t\, min\left(1;\ \left(\frac{\frac{Y^d{}_t}{Y^s{}_t}}{(Y^d/Y^s)^{noAI}{}_t}\right)^{\nu}\right)$$

In the fixed-point relationship, c^y and c^{tr} represent, respectively, the component of consumption proportional to output and the component linked to transfers, obtained by aggregating the household relationships. This explanation makes explicit the meaning of the coefficients, which the equations use in compact form. They are not new propensities to consume estimated from data. Y^d in the demand relationship denotes aggregate demand, while Y^dg in the household relationship denotes a group's disposable income.

Reading the Formula One Step at a Time

The fixed point takes into account the consumption generated by the income from output itself. The denominator contains one minus the component that returns to consumption: values are needed for which the division is defined and the system makes economic sense. The demand-to-capacity ratio is compared with that without AI. min with one prevents this rule from pushing output above potential; ν governs the response to the relative shortfall.

An Example to Get Your Bearings

Imagine a teaching relationship in which demand is twenty plus half of output. If output and demand coincide, fifty does not close the account: demand would be forty-five. Forty closes it, because twenty plus half of forty is forty. It is an example of a fixed point, not a calibration of the simulation.

The Limit to Remember

The model assumes that the no-AI path operates at potential. The result depends on this closure, not only on the technological savings.

A7 The Assumptions That Matter in Sensitivity

Chapters 9 and 16.

A variance measures the dispersion of results around their mean. The index compares dispersions, not numbers of people. No knowledge of the economic equations is needed to follow this section.

How much does the result change when a particular assumption changes?

The Mechanism and the Original Relationships

The technical relationship accompanying the reading is:

$$S^{T}_{i} = \frac{E[Var(Z \mid X_{-i})]}{Var(Z)}$$

Z is the outcome analyzed; X_i the assumption being moved; X_{-i} the set of the others. In the numerator, X_i is allowed to vary while the other assumptions are held fixed, and the variance obtained in this way is averaged. The denominator is the overall variance. The total index includes the parameter's own contribution and its interactions with the other levers.

Reading the Formula One Step at a Time

Var measures how much the results are dispersed. The vertical bar means that the other assumptions are held fixed. E here denotes an average, not exposure or the employed: the variation of the lever is repeated in many configurations of the others and av-

eraged. Dividing by the overall dispersion gives an index with no unit.

An Example to Get Your Bearings

If, in an abstract example, the average variability associated with a lever and its interactions is two and the overall variability is ten, the total index is 0.2. It does not mean that the lever concerns twenty percent of workers.

The Limit to Remember

Total indices include interactions and should not be added up as separate shares. Intervals and the precision of the estimator affect the reading.

A8 Public Revenue and Debt

Chapters 11 and 20.

Y is output, s^{L} the labor income share and Cg the consumption of group g. τL, τK and τC are hypothetical levy parameters on labor, capital and consumption for unidentified economies: they do not describe tax rates in force. G denotes government purchases; P pensions, not population; D debt, not labor demand.

How do we move from one year's balance to the following year's debt?

The Mechanism and the Original Relationships

The account includes taxes on labor, capital and consumption, government purchases, pensions, benefits and additional interventions. The equations distinguish the year's balance from the debt accumulated before that year.

P is pension spending, UB unemployment benefits, r the real interest rate and PB the primary balance. X collects policy spending, including redistributed transfers and training. The primary balance

excludes interest, which appears in the accumulation relationship for debt D.

$$PB_t = \tau_L s^L_t Y_t + \tau_K(1 - s^L_t)Y_t$$

$$+ \tau_C \Sigma_g C_{g,t} - G_t - P_t - UB_t - X_t$$

$$D_t = (1 + r)D_{t-1} - PB_t$$

Reading the Formula One Step at a Time

In the first relationship the tax rates apply to their respective bases; revenue is positive and spending is subtracted. In the second, one plus the interest rate multiplies the previous debt. Subtracting a negative balance is equivalent to adding a deficit.

An Example to Get Your Bearings

With initial debt of one hundred, interest of three percent and a primary surplus of two, final debt is one hundred and one. With a primary deficit of two, it reaches one hundred and five instead. These are teaching amounts in the same unit; the debt-to-GDP ratio also requires final output.

The Limit to Remember

GDP growth is not enough to guarantee a fiscal improvement. Distribution, tax bases, tax rates and benefits all matter.

A9 Checking the Counterparts

Chapters 11 and 23.

Y is output; C, I and G are consumption, investment and government purchases. PB is the primary balance, rD interest on debt and X^G policy spending included in demand. A5 and A8 explain how household and government accounts are formed.

Are the accounts of the different actors compatible with each other?

The Mechanism and the Original Relationships

The identity is checked in every period. Sh denotes household saving; Пret retained profits; I investment; PB minus interest the overall public balance. Payments abroad and payments received complete the regional balance. The right-hand side measures output minus aggregate demand, including XG, the component of policy spending that enters demand.

$$S^{h}_{t} + (\Pi^{ret}_{t} - I_{t}) + (PB_{t} - rD_{t-1})$$

$$+ (paid\ abroad_{t} - received_{t})$$

$$= Y_{t} - (C_{t} + I_{t} + G_{t} + X^{G}_{t})$$

Reading the Formula One Step at a Time

Read each bracket as a balance: households, firms, government and foreign payments. On the other side, output and the components of demand are compared within the scope implemented. It is not required that every balance be zero in individual regions: it is the links between accounts that must respect the relationship.

An Example to Get Your Bearings

If a household receives ten and spends eight, it saves two. Those two must find a counterpart in the other accounts: they cannot be counted as consumption and saving at the same time. The example clarifies why the check is needed, without replacing the full identity.

The Limit to Remember

Accounting consistency does not show that the rules of economic behavior describe the future.

A10 The Arrows of the Task Model

Chapter 23.

ΓL is the labor content in summary notation; E is exposure and a adoption. φ is the substitutable part, ψ the recovery per replaced unit, ℓ the lag and d_0 the baseline erosion. π denotes the cost saving and μ the gain in assisted tasks. A1–A4 offer the broader path.

How do the arrows of automation translate into calculable relationships?

The Mechanism and the Original Relationships

The two summary relationships form a compact formulation: the first follows the labor content of tasks and the second links its changes to growth. They keep their own notation, distinct from the full and revised relationships of Sections A3 and A4.

$$\Gamma L_{,t} = \Gamma L_{,t-1}\,(1 - d_0 - \delta_t^{AI} + \rho_t^{AI})$$

$$\delta_t^{AI} = \varphi\, \Delta(E_t\, a_t)$$

$$\rho_t^{AI} = \psi\, \varphi\, \Delta(E_{t-1}\, a_{t-1})$$

The three lines form the first relationship. ΓL represents the labor content of production; t denotes the year. The previous content is reduced by baseline erosion d_0 and by the substitution attributed to AI δ, while reinstatement ϱ supports it. The model assumes unit elasticity of substitution between tasks: in this specification the labor share coincides with the labor content of tasks. It is an assumption of the model, not an identity valid for every possible technology.

The product E_t a_t combines exposure and adoption. E here is exposure to AI, a adoption; they are not the employed of the age module. Δ indicates how much that product changes relative to the previous period. φ is the substitutable fraction of exposure. Annual substitution therefore depends on the increase in actual adoption of the exposed part, not on repeatedly applying the whole exposure to all workers every year.

Global exposure of forty percent does not mean that forty workers out of a hundred must be eliminated every year. It is a measure of technical possibility; annual substitution in the equations concerns instead the increase in adopted exposure.

The third relationship contained ℓ, the lag. Reinstatement in the current year looks at the exposure adopted in an earlier year; ψ indicates how much labor is reinstated relative to past substitution. The process therefore has a memory. For some years substitution can advance before reinstatement catches up. Changing ℓ or ψ changes the path of the gap, even if output grows.

Training could help reduce some lags, but the name of a course cannot replace a model parameter. Required skills, available jobs and access must be measured.

d_0 keeps, in the world without AI, the erosion not attributed to the new technological wave. The paired simulations use the same baseline parameters: they isolate a difference within the model. A causal attribution in the data requires instead a credible empirical comparison.

Producing More Does Not Complete the Account

The second relationship adds up three contributions to growth due to AI. The three lines belong to the same equation and must be read together.

$$g_t^{AI} = \Gamma L_{,t-1}\, \delta_t^{AI}\, \pi$$

$$+ \mu(1 - \varphi)\, \Delta(E_t\, a_t)\, \Gamma L_{,t-1}$$

$$+ \frac{1}{2}\, \pi\, \Gamma L_{,t-1}\, \rho_t^{AI}$$

The first contribution combines previous labor content, substitution and the cost saving π on automated tasks. The second concerns assistance to non-replaced tasks: μ expresses the gain of labor augmented by technology and the factor one minus φ distinguishes the non-substitutable part. The third attributes a contribution to the reinstatement of new tasks, with the coefficient of one half set in the specification.

The coefficient of one half is part of the productivity rule adopted, not a new empirical measure. Like the parameter ranges, it requires discussion and testing. The three terms specify the link between adoption, substitution, assistance, new tasks and growth, without by themselves determining which household receives the benefit.

Reading the Formula One Step at a Time

Δ means change relative to the previous period. E times a joins technical possibility and actual use. The recovery term looks back ℓ years; ψ governs its size. AI growth is a sum of three contributions, not three rates to be applied one after another. For the economic meaning of potential output, see A2.

An Example to Get Your Bearings

If adopted exposure goes from 0.10 to 0.14, its increase is 0.04. With a substitutable fraction of 0.5, the substitution term is 0.02. A lagged recovery does not necessarily appear in the same year.

The numbers are illustrative and do not describe an observed activity.

The Limit to Remember

This summary formulation of the model keeps its own notation and must be distinguished from the revised wage relationship in A4.

A11 Sharing Working Hours

Chapters 24 and 25.

h is the remaining share of individual working hours. b is the budget as a share of GDP, r the part of the lost wage that is made up and s^L the labor income share. r here is not the interest rate used in A8.

How much work can be shared with a given budget?

The Mechanism and the Original Relationships

$$h = 1 - \frac{b}{r \times s^L}$$

h is the remaining fraction of hours, b the share of GDP devoted to the measure, r the share of the lost wage that is made up, s^L the labor share. Capacity in people grows by a factor of 1/h. The assumption allows divisible work to be shared: shifts, skills, fixed costs and coordination could make it harder to apply in a firm.

Reading the Formula One Step at a Time

The remaining fraction of hours is obtained by subtracting from one the ratio between the budget and the cost of the top-up. The factor one divided by h translates the same quantity of hours into people. h must be positive and no greater than one, and the denominator must be defined.

An Example to Get Your Bearings

If in the example h is 0.8, each person keeps eighty percent of their hours. One hundred full-time units of labor correspond arithmetically to one hundred and twenty-five people at those hours, because one hundred divided by 0.8 is one hundred and twenty-five.

The Limit to Remember

The arithmetic step does not show that a firm can actually distribute all tasks in that way.

A12 Concentration and the Aggregate Labor Share

Chapter 27.

HHI is the index of sales concentration; s_i is the market share of firm i. VA_i is its value added and λ_i the share of that value attributed to labor. The shares s_i can use the 0–1 or 0–100 scale, as long as it is consistent. No preliminary section is needed.

How can concentration be summarized without losing track of the measure used?

The Mechanism and the Original Relationships

$$HHI = \Sigma_i s_i^2$$

The aggregate labor share is an average of firm shares weighted by value added. With shares of sixty and thirty percent and equal weights, the result is forty-five percent. The following formula makes those weights explicit.

$$Labor\ share = \frac{\Sigma_i VA_i \lambda_i}{\Sigma_i VA_i}$$

VA denotes the value added of firm i and λ its share attributed to labor. The sum in the numerator reconstructs labor income, the

one in the denominator total value added. This teaching calculation assumes positive values and consistent definitions. It is not a new equation of the software.

Reading the Formula One Step at a Time

Σ means adding up one term for each firm; squaring gives more weight to large shares. For the aggregate labor share, by contrast, the firms' labor incomes are added up and divided by total value added. Firms therefore weigh according to value added, not one each.

An Example to Get Your Bearings

With four firms at twenty-five percent, the sum of four squares of twenty-five is 2,500. With shares of seventy, ten, ten and ten it becomes 5,200. With shares between zero and one, the same results are 0.25 and 0.52.

The Limit to Remember

Concentration alone does not prove abuse or effects on prices. The aggregate share requires consistent scopes and definitions.

Symbols and Definitions

GDP. Value added produced in a territory over a period. GDP growth does not indicate how income is distributed.

Percentage points. Difference between two percentages: from 10% to 12% is two percentage points, equal to relative growth of 20%.

Median and percentiles. The median divides the ordered values in half. The 5th and 95th percentiles bound the central 90% of the model's draws; they are not automatically a confidence interval.

Paired simulation. Two runs with the same parameters, one with the AI wave and one without. It isolates the effect within the model; it does not prove causality in observed data.

Labor supply and demand. People available to work and labor required by production. The potential supply of the age module includes states that do not coincide with the official labor force.

Model symbols. t: year; r: region or, in the debt formula, interest; g: income group; Γ: labor content; δ: substitution; ϱ: task creation; ε: wage elasticity; Y: output; N: macro employment; PB: primary balance; D: debt. In the age module E denotes the employed, while in the macro equations E denotes exposure to AI.

The Code Behind the Book

The source code of the software used for the simulations in this book can be downloaded free of charge from the repository:

https://github.com/0xENX/aimodel_replication

It contains the code, the observed data and the results cited in the chapters, and makes it possible to rerun the calculations.