

WORK AFTER AI.

The annual assessment of how artificial intelligence is transforming work

The Great Misconception About the AI Job Killer

Contents

01	Foreword	3
02	The Ten Key Findings of 2026	4
03	What a trend is and what it is not, and what this report does not know	6
04	Method and disclosure	7
05	Three forces in the labor market, one wrong headline, and the question that really counts	10
06	The AI argument, the loudest of the three forces, says more about the capital markets than about the technology	14
07	The career gap is measurable, but its cause is not so easy to pin down	25
08	The second axis of the career gap runs along gender, in risk as much as in access	41
09	The long rebuild is the most consequential of the three forces, and for now it is failing not because of the machine but because of the organization	44
10	The technological foundation is being built faster than companies adopt it into their work	64
11	The organization of the future arranges its work around topics instead of functions, and it gets smaller in the process	69
12	Anyone who transfers the American debate to Europe is measuring with the wrong yardstick	75
13	The real European controversy is being fought out in the individual countries, not in the aggregate figures of the union	78
14	Dismissal protection is the structural reason why the same shock lands more quietly in Europe	81
15	Demographic pressure inverts the whole starting question	83
16	European sovereignty is a strategic lever, not a compliance argument	87
17	Nine predictions for 2027 that this report will grade publicly next year	93
18	The future of work is decided not by the machine but by the company's ability to work with it	97
19	Sources	103

Foreword

This report grew out of a contradiction. In the headlines, artificial intelligence is eating jobs. In the official statistics, you look for the job killer in vain. What you find there instead is something else, and it is the same in almost every data series: a technology that is being tried out widely and is only rarely really productive in use, an entry level that is changing without anyone being able to name the cause precisely, and an organization that does not yet know how it is supposed to work with a colleague who is not a person.

The title of this first edition is a claim, and the report is its test. It takes the job killer headline apart into three processes that have different causes and run at different speeds: the AI argument, the career gap and the long rebuild. It tests each of them against the best available evidence, against official statistics, against peer reviewed studies, against the vendors' own usage data and against documented cases from companies. Where the findings contradict each other, the text says so. Where a piece of evidence comes from someone with an interest of their own in the result, that is noted alongside it. Where a number could not be verified, it does not appear.

In 2002, the economist Carlota Perez described in "Technological Revolutions and Financial Capital" how every great technological revolution since the industrial revolution falls into two halves. In the first, the installation phase, financial capital drives the new technology forward with exuberance and bubbles; in the second, the deployment phase, the gains only arrive across the board once institutions, ways of working and skills have adapted to the technology.¹ Lay this pattern over the year 2026 and almost all of it is recognizable: the investment wave, the valuations, the loud announcements, and behind them an economy that has not yet taken the technology into its work at all. The job killer belongs to the first half of this story. The questions this report is about belong to the second. How quickly does an organization build the ability to work with the technology? And what happens to those that do not?

From now on, Work After AI will appear every year in the autumn, in the same structure, with the same core questions and with a prediction chapter that I grade publicly in the following year, my own predictions as well as those of others. The real value is not meant to lie in the single edition but in the series, which over the years makes visible what has actually changed and what was only announced. Part of that is disclosing my own interests, because I earn my money with enterprise software that uses artificial intelligence. The report is therefore explicitly not tied to any vendor, any model and any consulting firm, and it recommends nothing that could be bought. It orders the discussion. What consequences to draw from it is for each company to decide for itself, and that is exactly the point.

Gerhard Kürner, Wels, September 2026

The Ten Key Findings of 2026

- 1 **The job killer is a capital markets story.** Only about 5 percent of American companies that use artificial intelligence report any effect at all on their headcount, and the technology replaces software a good three times as often as it replaces people. The job cuts that are justified with it follow the logic of the stock market, not of the companies.²
- 2 **Expectation outruns measurement.** In 2025, 32 percent of companies expected their workforce to shrink through AI in the following year, 14 percent actually experienced it a year later, and 39 percent nevertheless now expect it for next year. The same companies missed their prediction and then expect more of it.³
- 3 **The career gap arises in hiring, not in layoffs.** Employment among 22 to 25 year olds in the most exposed occupations in the United States is 19 percent below the level of their peers in less exposed occupations, and the entire gap is explained by declining hiring. Whether the machine, interest rates or remote work is the cause is open; in Europe the drop is so far measurable only in Switzerland, which can mean protection or delay.⁴
- 4 **The entry level is shifting from execution to oversight.** When the machine takes over the routine work that used to be the learning task of the entry level, the entry level has to be rebuilt. Dual vocational training, which has never separated learning from working, is the underrated lever for this in the German speaking countries.⁵
- 5 **Women carry the higher risk and have the lesser access.** In Germany, men use generative AI 20 percent more often than women in private life and 29 percent more often at work. Among 16 to 24 year olds, women are ahead. The gender gap arises in the workplace, and it closes where the employer pays for the training.⁶
- 6 **What the individual gains does not arrive in the organization.** 80 percent of respondents report productivity gains of their own. 37 percent see any value at all in the operating result, but only about 6 percent a measurable contribution to operating profit of at least 5 percent. Of those 6 percent, just under three quarters rebuilt their workflows from the ground up. The long rebuild does not fail because of the machine, it fails because of the organization.³
- 7 **Delegating is a leadership decision, and whoever stops checking loses the skill.** The willingness to hand work over to AI follows only to one third from the assessment of the technology. At the same time, a company loses capabilities it has long possessed as soon as the machine's output is adopted unchecked; in a field experiment, consultants working with AI outside the machine's performance frontier were wrong 19 percentage points more often than without it.^{7, 8, 9}

- 8 The organization of the future orders its work by topic instead of by function, and in doing so it becomes smaller.** 43.5 percent of work related AI use concerns tasks that until now belonged to a different occupation; small teams take on adjacent specialist and process topics as well, and the specialist department becomes a review body. The HR function itself is caught in the pilot trap: 28 percent of its processes run productively with AI, its own expectation for the share of automatable tasks has fallen within a year from over 30 to 20 percent, and employees report 3.4 training days a year while the HR departments estimate 6.2.^{10, 11, 12}
- 9 Marketing is the first function in which the rebuild becomes visible, from the inside as well as from the outside.** From the inside, sales reaches into online marketing with agent supported teams and management reaches into brand management, while marketing itself does not reach out. From the outside, retail media grows to around 175 billion dollars and overtakes television in 2026, 82 percent of the members of the American association of advertisers run an in-house agency of their own, and the range among the agency holding companies runs from WPP at minus 5.4 percent to Publicis at plus 5.6 percent. The crisis hits the model of selling time, not the category.^{13, 14, 15, 16}
- 10 In Europe, demographic pressure reverses the opening question, and the future of work is decided by the company's ability to work with the technology.** By 2040 Germany will lose around 13.3 million people in the labor force to retirement, that is 30 percent, and the younger cohorts do not make up for it. In the German speaking countries, artificial intelligence is the most likely instrument for holding on to prosperity, provided companies take it into their work before the experienced ones retire with their knowledge.¹⁷

What a trend is and what it is not, and what this report does not know

This report distinguishes three things that constantly run into one another in the discussion about artificial intelligence and work: what is measured, what is assessed and what nobody can know. Measured are the three forces around which the main part turns. The AI argument can be shown in layoff statistics, company surveys and capital market communication, the career gap in payroll and register data, the long rebuild in surveys, official statistics and by now in the usage data of the model vendors themselves. These three forces are trends in the strict sense, because they have a direction that appears in several independent sources and over more than one year, and because it can be checked in the coming year whether they hold. Assessed are the predictions and the outlooks at the end of the chapters; they state the author's view, are derived from the documented findings and are marked in the text as assessments. Nobody can know how fast the technology itself gets better. This report makes no statement about when a model reaches which capability. It cites the scenarios of others as benchmarks and grades them, but it does not place a scenario of its own beside them. Its central thesis does not depend on that abstention, because it holds that the rebuild of work is for now not failing because of the machine but because of the organization, and that applies to every speed of technical progress that is plausible in 2026.

What this report does not know belongs on this page as well. It does not know whether it means protection or delay that no effect has so far shown up at the entry level in Europe, because the spread of the technology here lags the American one by two to three years, and the measurements that could settle this question will not exist before 2027. It does not know whether the gap between pilot and production is a transitional phenomenon that closes as organizations learn, or whether, as the Notion survey suggests, it opens further with growing maturity. It does not know how the effect looks in firms with fewer than ten employees, because there is no reliable survey for that in Europe, and it knows next to nothing about wages, training and codetermination in small firms. It does not know what share of the American layoffs the technology actually accounts for, because the statistics count the reason the company states, not the one that applies. And it does not know whether the usage data of the model vendors, which it treats in several places as the sharpest measurement so far, will still count as such in five years, because none of these data series can be recalculated from outside. These gaps are not weaknesses of the report that could have been avoided, but the state of knowledge in September 2026, and they are the reason the report appears annually.

Method and disclosure

The report places other people's evidence in context and makes assessments of its own, and it keeps the two visibly apart. The evidence base of the founding edition consists of the sources listed in the sources section, among them official statistics from Germany, Austria, Switzerland, France, the United States and the European Union, peer reviewed studies, working papers from central banks and research institutes, the usage data of the three large model vendors, surveys by consulting firms and software vendors, and documented company cases. The same procedure applies to every source. It is read in the original, and every number used in the text is checked against the primary reference and, where a second independent source exists, cross checked against that one. Divergences between the two are stated in the text or in the sources section and are not smoothed over, and numbers that failed this check do not appear in the text, even when they are circulating in the press. For every source there is an internal dossier with its verification status, its limitations and a list of blocked values, and that dossier is the basis for future editions. Commercial senders and senders with an interest, that is consultancies that sell the rebuild they describe, vendors whose business depends on the spread of the technology, associations and companies speaking for themselves, are marked as such in every place; their numbers are carried as self reported statements, never as measurements. Job postings are never read as an employment effect, and an occupation's exposure to the technology is never equated with its replacement, because those two confusions are the most common errors in the public debate.

The report relies exclusively on sources that every reader can check for themselves. Up to version 1.1, assessments from the author's working conversations with owners and management were also included. They were dropped with version 1.2, because these conversations were neither conducted according to a sampling plan nor recorded and therefore cannot be verified from the outside. Since this version, the report makes assessments of its own in only two marked places, in the predictions and in the "Looking ahead" section at the end of the chapters, and both are derived from the documented findings. From the 2027 edition on, the DACH board pulse is added, a fixed catalog of core questions asked in the same form every year; the questionnaire, the number of participants and the field period are published with every edition. Added to it then is the report's lead metric, the productive rate, that is the share of companies in which at least one AI system handles a core task in regular operation, outside of pilot, test and one off initiative. The founding edition does without both, because a measurement that exists only once is not a time series.

Artificial intelligence worked on this report, and that is stated here. Research, source checking, drafts, editing and the management of the source inventory ran on the platform of 506.ai, in a working environment in which every check is documented; for the agentic checking of sources, that is the independent locating of the primary reference, the comparison of every number against a second source and the keeping of the blocklists, dedicated skills were developed that run this procedure in the same form for every source. The theses, the selection of sources, the predictions, every decision to include a finding and every sign off come from the author, and every text has been read back by him. The report does not carry this way of working as a footnote but as part of

its subject, because it describes a division of labor between human and machine that it practices itself, with the same limits it describes in the chapters. The machine invented numbers several times where a figure contained none, it attributed sources wrongly and confused versions, and each of these errors was found only because a human compared the statement against the primary source.

The report is self funded, and its results are not for sale. It is paid for by no vendor, no consultancy, no association and no public body, it carries no advertising, and none of the organizations cited saw the text before publication. At the same time, the author discloses his interests. Gerhard Kürner is CEO of 506.ai, the European platform for service as a software and agentic engineering, and holds further roles in the field of enterprise software. None of these companies is a source, a piece of evidence or a case example in the report, the report is tied to no model and no vendor and recommends no product. The publisher is Kuerner Group GmbH, Dragonerstraße 54, 4600 Wels, Austria; the author is Gerhard Kürner. The present edition is registered under the DOI 10.5281/zenodo.22868155, the author under the ORCID 0009-0003-0611-0526. The report is cited as: Kürner, Gerhard (2026): Work After AI 2026. The Great Misconception About the AI Job Killer. Work After AI, 2026 edition, version 1.2. Wels: Kuerner Group GmbH. DOI 10.5281/zenodo.22868155. The report is published under the Creative Commons Attribution 4.0 International license (CC BY 4.0), which permits redistribution, quotation and adaptation with attribution of the author, and every edition remains reachable at a permanent address. The present version carries the version number 1.2 and was published on September 24, 2026; the first publication, version 1.0, appeared on September 16, 2026, and version 1.1 on September 18, 2026. The 2026 edition will be updated through the end of the year when new data sharpen the findings or errors need correcting. Every version carries its own version number with a date. The current version is at workafterai.org, and what has changed between versions is set out in the version history at the end of the report. From version 1.2 on, every version also appears as a machine edition, that is as a text edition with a fixed identifier for every chapter, key finding, prediction, external benchmark and source, as data files and as a skill package for assistants; the files sit next to the PDF under the work DOI and are described at workafterai.org.

The report works on three levels that move at different speeds. The edition is the annual level. It appears in the fall, and it is the only place where the previous year's predictions are graded and the measurement series are extended by one year. The version is the quarterly level. Within an edition, up to three further versions follow at intervals of roughly one quarter, each with a date, its own version number and its own permanent identifier (DOI), so that every citation refers to a fixed state. A version appends new sources and sharpens key findings; the predictions are neither changed nor withdrawn within an edition, because a prediction that is adjusted along the way can no longer be graded at the end of the year. A version is published only when new data change a finding or an error needs correcting; a version history without entries does not justify a version. The third level is the analyses on workafterai.org. They appear between versions whenever a development calls for it rather than on a calendar, and each analysis names the key finding it touches and says whether it confirms it, shifts it or leaves it open; it is incorporated with the next version. This rhythm follows the rhythm of the sources, because the official statistics and the large surveys the report relies on appear annually or quarterly, and a report cannot have new findings more often than its sources.

For the founding edition 2026, published in September, an exception applies: version 1.3 is planned for early November after the fall studies, and the final version 1.4 for December 31, 2026.

The report develops further, and it states every change to its method. The structure around the three forces named above, the chapter on the organization of the future, the prediction chapter with its block of external benchmarks and the rules of source work stay the same from edition to edition, so that a longitudinal view emerges. From 2027 on, the board pulse, the productive rate, the impact matrix, which sorts the effect of the five changes to work, that is the AI argument, the career gap, the long rebuild, the new structure of work and demographic pressure, across 14 industries by time horizon, and the public grading of the previous year's predictions, our own as well as those of others, are added. What is an assessment in this edition is meant to be a measurement then, and where a definition changes, that is recorded in the method section of the edition concerned, so that the years can be compared.

Three forces in the labor market, one wrong headline, and the question that really counts

The headline that tries to explain too much at once

Hardly any story has taken hold as effortlessly over the past year as the one about artificial intelligence pushing people out of work. It has the advantage that it appears to confirm itself the moment you open the news. One large technology company announces thousands of job cuts on a single day, another eliminates entire departments, and in almost every one of these statements the word appears that promises to give the whole thing its meaning. The technology takes over, the human steps aside, and the layoff statistics deliver the evidence along with it.

The story is catchy, and in this form it is wrong. Not because artificial intelligence leaves no traces in the labor market, but because the headline fuses three completely different processes, which in truth rest on different causes and run on different time scales, into a single story. Anyone who takes it apart sees a picture that is less dramatic but far more consequential than the one in the headline. And they see where the question that actually matters lies, the one that is almost entirely lost in the excitement about the machine.

“The headline about AI the job killer fuses three processes, which in truth rest on different causes and run on different time scales, into a single story.”

Three processes with their own pace sit inside one headline.




FIGURE 1

The first force is the AI argument, a wave of layoffs that calls itself artificial intelligence and goes by another name

The most visible of the three forces is the wave of job cuts, and it is at the same time the one that has the least to do with the technology in whose name it happens. A survey by the consulting firm Gartner (gartner.com) among 350 large companies, all of which already work with artificial intelligence, did find that four out of five of these companies have cut staff. But it also showed that there is no connection between the depth of the cuts and the economic success of the investments. Those who cut the most did not end up better off, and in some cases worse. That is a remarkable finding, because it contradicts the internal logic of the headline. If the job cuts were really the result of a more productive machine, they would have to show up in the numbers, and that is exactly what they do not do.¹⁸

What shows up instead is an older and more sober pattern. A substantial part of what is presented today as technological modernization is the correction of overstaffing that built up in the years of cheap money. Investors from Silicon Valley, who are hardly suspected of underestimating the importance of the technology, say openly that nearly every large company is overstaffed, many by a quarter, some by half and more.¹⁹ Artificial intelligence supplies the better argument for these cuts, because a reduction presented as modernization lands differently in the capital markets than a plain cost saving decision. This first force is therefore real in its consequences for the people

affected, but it is wrongly named in its cause. This report therefore gives it the name that describes its mechanics, the AI argument.

The second force, the career gap, hits those starting out, and its cause is disputed in the research

The second force is subtler and more serious. It shows up not in the headlines about individual corporations, but in the data on young people entering working life. A widely noted study by the Stanford Digital Economy Lab showed, on the basis of actual payroll data, that employment among 22 to 25 year olds in those occupations that are particularly exposed to artificial intelligence has declined markedly in comparison, while older and more experienced employees in the same occupations remained stable or even gained.⁴ Here, then, the technology appears to be doing exactly what the headline claims.

But at this point the matter becomes harder rather than simpler, because the research is anything but united about the cause. A study from the Federal Reserve by Jessica Liu and Douglas Webber covering more than a million companies finds no measurable effect,²⁰ and an analysis by the payments provider Ramp (ramp.com) together with the labor market data provider Revelio Labs (reveliolabs.com) does confirm the collapse at the entry level, but shows that it is precisely not driven by those firms that are actually adopting artificial intelligence.²¹ The collapse is therefore demonstrable, but attributing it to the technology is contested, and a substantial part of it is likely to go back to the sharpest turn in interest rates in four decades, which made every new hire at the entry level more expensive overnight. Anyone who wants to treat this second force honestly has to endure the dispute instead of resolving it in favor of the simple story. Because its finding appears in the same place regardless of how that dispute ends, namely at the beginning of the working life, this report carries it under the name of the career gap.

The third force, the long rebuild, is the actual effect of the technology, and it is slower than its reputation

The third force is the real change that artificial intelligence sets off inside companies, and it has the least in common with the pace of the headlines. It does not run through the spectacular replacement of entire workforces, but through the slow and laborious reworking of processes, through the question of which task a digital team member can take over and which it cannot, and through the far harder question of how an organization builds the necessary knowledge in the first place. This force is the most consequential of the three, but it is invisible, because it cannot be condensed into a press release. It happens in the engine rooms of companies, not on their front pages. This report calls it the long rebuild, because what decides success is not the new machine but the patient reorganization of work around the machine.

Why the three forces are constantly confused

The fusing of the three into a single story is not ill will but an understandable confusion, because all three are happening at the same time. The AI argument, the career gap and the long rebuild fall into the same window of time, and the human need to link simultaneous events causally does the rest. Artificial intelligence is at the center of attention anyway, and so it becomes the obvious explanation

for everything that is moving in the labor market at once. The analytical work does not consist in denying one of these three forces, but in separating them cleanly from one another and showing the real mechanics of each.

This separation is the subject of the present study. It rests on the publicly available quantitative research, which it reports honestly in its contradictions instead of picking out the convenient findings. It places an emphasis on the European labor market, which is a blind spot in the largely American discussion, even though its structures differ considerably. And it uses only numbers that have been checked against their primary source and, where one exists, against a second independent source. Only with this separation and this checking does it become possible to see behind the headline.

What this is really about in the end

Once the three forces have been separated, the question that really counts shifts to an unexpected place. It is then no longer whether artificial intelligence destroys jobs, because that question, as will become clear, is posed wrongly. It is how fast an organization builds the knowledge that makes the technology productive in the first place, and whether, in the haste of cutting jobs, it is destroying precisely the capability that will later decide its success. This question stands at the end of the study, not at its beginning. The path there leads through the three forces, and it begins with the one that is the loudest and delivers the least of what its name promises.

The AI argument, the loudest of the three forces, says more about the capital markets than about the technology


The wave of job cuts is the most visible of the three phenomena, and it supplies the seemingly most crushing evidence for the thesis of the machine as job eater. When a corporation cuts thousands of jobs in a single day and speaks of artificial intelligence in the same breath, the connection seems obvious. On closer inspection, however, it dissolves, and a more sober explanation takes its place, one that has more to do with the peculiarities of the capital markets than with the capability of the technology.

The sharpest evidence for this comes from a study by the consulting firm Gartner from the spring of 2026, which covered 350 large companies each with more than a billion dollars in revenue, every one of which was already working with artificial intelligence. Four out of five of these companies had cut staff. The decisive finding, however, does not lie in that rate but in what shows up between the depth of the cuts and their economic success, namely nothing. There is no correlation between the extent of the cuts and the success of the AI investments, and those who cut the most were not better off in the end, in some cases worse. Helen Poitevin, who is responsible for this research at Gartner, sums it up in the sentence that many chief executives reach for layoffs in order to show quick AI wins, but that this attitude is misguided, because job cuts create budget room and not success. Part of the context is that Gartner has not made this survey available as a standalone publication; what is public is only the analyst's commentary and the figures quoted from it, while the methodology and the raw data are not disclosed, and the finding is therefore to be read as a consulting firm's survey whose results a reader cannot recalculate.¹⁸

That this pattern is now also spoken aloud where the labor market data originate is shown by a remark by Virginia Sondergeld, economist at the Hiring Lab of the job board Indeed (indeed.com), in Die Zeit in September 2026: Companies, she said, prefer to justify job cuts to investors and the public with artificial intelligence rather than with the weak economy. The sentence appeared in an article about the collapse of entry-level positions in Germany whose headline adopted exactly this justification, while the data beneath it date the beginning of the decline to the spring of 2022, half a year before ChatGPT.¹⁴¹

So far companies use AI to replace software, hardly any staff.

Share of AI-using companies in the United States, in percent



Source: U.S. Census Bureau, Business Trends and Outlook Survey, Nov. 2025 to Feb. 2026 (Bonney et al., 2026). Self-reported.

FIGURE 2

This finding contradicts the internal logic of the headline. If the job cuts really were the result of a more productive machine, they would have to show up in better numbers. That they do not points to a different mechanism, and this mechanism is neither new nor mysterious. A large part of what today goes under the name of technological modernization is the belated correction of overstaffing from the years of cheap money. Venture capitalists, whom nobody suspects of playing down the importance of the technology, name the true causes without varnish. Marc Andreessen, for example, considers practically every large company overstaffed, many by a quarter, most by half, some by three quarters, and traces this back to the zero interest rate period and the frantic build-up of workforces during the pandemic. Even Sam Altman, whose business rests on the spread of AI, concedes that there is a certain amount of cheating going on and that firms attribute layoffs to artificial intelligence that they would have carried out anyway, while deliberately holding back on the exact share.¹⁹

But why do so many companies choose this argument of all things when the real reason is a plain decision to save money? The answer lies in the reaction of the capital markets. Since the first broadly available language models appeared at the end of 2022, according to an analysis by J.P. Morgan (jpmorgan.com), around three quarters of the price gains in the broad American stock index have come from the field of artificial intelligence.²² In such a climate, a cut presented as AI modernization lands with investors in an entirely different way than a sober reference to costs that are too high. At the same time, personnel costs are one of the largest freely controllable blocks of spending, so that reducing them cross-subsidizes the high investments in the technology. Job cuts finance AI investments without measurable success, and AI supplies the argument that justifies the job cuts.

“Job cuts finance AI investments without measurable success, and AI supplies the argument that justifies the job cuts.”

At this point caution is called for, because the trackers circulating online that record layoffs day by day are no use as evidence. They throw job cuts from the technology industry into one pot with those from the financial industry and from health care, which need not have anything to do with the technology, and in this way they arrive at totals that do not stand up to a clean recalculation. What carries weight is therefore not the statistics of the daily reports but the survey by the consulting firm Gartner among 350 large companies cited at the outset, which robustly shows that the depth of the cuts reveals nothing about the benefit of the technology.¹⁸ The much cited example of a corporation that cut thousands of jobs despite brilliant quarterly figures also demonstrates the logic of the capital market rather than the effect of the technology, because a company that cuts from a position of strength rarely does so out of technological necessity.

How this mechanism can be observed in an individual case is shown with almost textbook clarity by the British neobank Starling Bank (starlingbank.com) in July 2026. The company announced the cutting of around 130 jobs, about 3 percent of its more than 4,000 employees, and publicly presented this as a simplification of the structure in the course of its AI build-out. The underlying business figures tell a different story. Revenue had fallen by 5.6 percent to 887 million pounds in the past financial year, pre-tax profit by around 3 percent to 217 million pounds, and Starling itself named as the main cause not its own technology but the lowered base rate of the Bank of England, which reduced interest income by 52.5 million pounds. The company's language, according to which structures were being simplified and duplication reduced, is the classic restructuring formula that existed long before generative AI, and nowhere does it say which of the eliminated activities are actually being replaced by an AI system. At the same time, Starling made clear that it would continue to hire AI and software developers, a pattern that describes a shift in the personnel mix rather than automation in the narrower sense. What is remarkable is that Starling can also point to a real, product-side AI commitment, from spending analysis to an agentic financial assistant, which is what makes the case so revealing, because what hides behind the argument here is not a company without AI substance but one with genuine AI substance, which additionally dresses a cost-cutting step forced by the interest rate environment in the story that sounds better on the stock market than the sober reference to falling interest income.²³

What can be inspected in miniature at a neobank with more than 4,000 employees was demonstrated in large format three days later by the largest software corporation in the world. Microsoft (microsoft.com) announced on July 6, 2026 the cutting of around 4,800 jobs, a little more than 2 percent of its roughly 220,000 employees worldwide, and the wire reports slotted the step straight into the AI layoff wave, in the German speaking countries even to the point of the formulation that the corporation was using productivity gains from artificial intelligence for the job cuts. The company itself said the opposite. Head of human resources Amy Coleman wrote explicitly in the memo to the workforce that the eliminated roles were not being replaced by artificial intelligence, though the technology was changing the way work gets done. The actual causes stood right next to it in the financial data, namely investments of 190 billion dollars in AI infrastructure in the current year alone, a share price that had lost around 30 percent and thus about 1.2 trillion dollars in market value within nine months, and the resulting pressure from the

capital markets on operating costs. The cross-subsidy logic of this chapter can hardly be illustrated more clearly, because in the same week in which the layoffs were issued, Microsoft launched a unit funded with 2.5 billion dollars that sends 6,000 developers directly into its customers' companies to make AI systems productive there. What is being replaced, then, is not people by the machine, but classic sales and consulting roles by more technical roles at the same point of value creation, while the job cuts help pay for the data centers.²⁴

The largest part of the cuts moreover falls on a business whose problems have nothing to do with the capability of the technology. The Xbox gaming division is losing sixteen hundred jobs in the first step and, by the end of the financial year, an expected 3,200, about a fifth of its global workforce, and the new division head Asha Sharma justified this in her memo with remarkable candor. The business was not healthy, it was working with margins three to ten times lower than those of comparable platform and publishing businesses, strategic bets such as the Game Pass subscription model had fallen short of expectations, and some work passed through up to 14 layers of management, which are now to be reduced to five at most. That is recognizably the vocabulary of a balance sheet repair and not that of automation. The only place where artificial intelligence actually appears in this justification is an indirect one, because the worst hardware crisis in the industry's history, which Sharma laments, is fed by chip prices that of all things the data center boom of the AI vendors is driving up. Where the technology acts on employment here, it acts through the capital and cost channel, not through the productivity of a digital colleague. That of all corporations the one that profits from the AI story like no other denies the replacement thesis internally and in writing, while the headlines ascribe it to that same corporation as a matter of course, makes the case perhaps the most revealing of the year so far.²⁴

That this mechanic is not an Anglo-Saxon peculiarity had already been demonstrated a few days earlier by Europe's largest software corporation, in a third variant that manages entirely without layoffs. On the evening of July 1, 2026, the SAP executive board announced to the workforce in an internal email that new hiring would in future be concentrated exclusively on selected profiles, above all on central AI roles that were decisive for long term success, that internal travel without an AI connection would be paused and spending on suppliers cut, all to finance what was described as a significant AI push. Bloomberg had seen the memo, a spokesperson expressly confirmed the thrust when asked, and customer facing work and critical AI projects would accordingly remain fully funded. The language policy alone is revealing, because while the international press wrote of a hiring freeze, the company itself spoke merely of greater discipline in hiring, spending and travel, and the German business press emphasized that this was not a general freeze but strategic steering. It is the same gap in interpretation as at Microsoft, only with the roles reversed, because there the press sharpened the AI attribution against the company's denial, whereas here the company softens the description of a cost brake that is harder than it may care to present it. In Walldorf too the causes lay not in the engine room but on the capital market. The SAP share had lost around a third of its value since the beginning of the year, because investors fear that artificial intelligence could hollow out demand for the classic software business, and one day before the memo the corporation had lost the bidding contest for the industrial AI firm Cognite, which went to Schneider Electric for 3.1 billion dollars. Here nobody is laid off, and yet the technology acts on employment, namely through the diversion of personnel and non-personnel budgets into AI budgets

under the pressure of a capital market that has already priced in the disruption of the corporation before it has taken place.²⁵

The case gains its real depth only against the background of its prehistory. SAP (sap.com) had completed a restructuring at the beginning of 2025 that had grown from an initially announced 8,000 to a gross figure of around 10,000 positions and cost 3.2 billion euros, while the total workforce remained largely stable thanks to simultaneous new hiring, a rebuild, then, and not a reduction. Chief financial officer Dominik Asam cast this logic into a remarkable formula in the autumn of 2025, when he announced that the corporation would in future permanently cut 1 to 2 percent of its workforce per year and compared this to brushing one's teeth. Restructuring is thus no longer an event but a permanent condition, a quiet, continuous reduction below the headline threshold that appears in no layoff statistic as a wave and yet takes effect year after year. Chief executive Christian Klein said in the spring of 2026 that it would be an illusion to believe that artificial intelligence drives productivity and the workforce nevertheless ends up looking the same, that fewer of certain roles that exist today would be needed. That of all corporations the one that was celebrated in the summer of 2025, thanks to its AI bet, as temporarily the most valuable company in Europe is a year later saving on hiring, travel and suppliers in order to keep financing that same bet makes it the European key case of this chapter and at the same time points to a question that the second part of this report explores in more depth, namely out of what depth of capital European companies are actually running their AI race.²⁵

The fourth case is supplied by the corporation that has advocated the replacement thesis most loudly, and it is the most revealing, because there the argument was not a pretext but the plan. According to a special report by the news agency Reuters from August 26, 2026, based on internal documents and conversations with more than 20 people involved, the leadership of Meta (meta.com) drew up in January 2026, under the name Project OT, the conversion into a company in which AI agents take over a large part of the daily work of thousands of employees and small cadres staffed with top talent supervise them. In the scenario planning, individual teams were shrunk by up to 60 percent, the cuts were to run in two waves in May and in November, and Meta has confirmed the plan, the two waves and the scenario to Reuters, but disputed the intention of cuts on this scale for the entire workforce. On the evening of May 19, a few hours before the first wave, Mark Zuckerberg called off the planning of the second. The first ran the next morning as scheduled and affected around 8,000 employees, that is 10 percent, and to the workforce Zuckerberg justified it not with the machine but with the capital expenditure that the corporation puts at 130 to 145 billion dollars for 2026. What led to the reversal Reuters could not establish; three circumstances that coincided are documented. The workforce openly pushed back, approval in the internal sentiment survey fell from 74 to 55 percent, and employees opposed the recording of keystrokes and mouse movements for training the company's own models with a petition. At the same time internal analyses that Reuters was able to see and that Meta has not commented on showed that the agent technology at the center of the plan did not deliver the hoped for gains: the number of changes to the internal platforms rose by 220 percent within a year, the number of changes that reached users as a new or improved feature by 36 percent, while serious technical and security incidents increased by 40 percent. On top of this came the investors' question of what the corporation had to show for its spending. At the beginning of July Zuckerberg conceded to the workforce that the technology had not developed as quickly as expected, and since then Meta has

publicly advertised that it empowers people instead of automating work; in his essay from August he writes that companies will become smaller, that there will then simply be more companies with fewer people.²⁶ The case is so revealing for this chapter because it shows the mechanic from the other side. At Starling, Microsoft and SAP, replacement was a story about a decision to save money. At Meta it was meant seriously, and it failed because of the technology and because of the people it was supposed to replace. What remains is what remains in all four cases: the cuts that finance the investments, and a story that adapts to the situation.

That these four cases are not isolated outliers but an industry wide pattern is also evident in how the corporate leaders themselves speak in public. On July 6, 2026, the Wall Street Journal described a remarkable shift in mood among the leading figures of the American technology industry, away from the grim prediction of mass job losses through AI systems and toward a markedly more optimistic tone. OpenAI chief Sam Altman conceded at a conference in May 2026 that they had largely been right about the technological predictions but pretty wrong about their social and economic consequences. Anthropic chief Dario Amodei, who in May 2025 had still warned of the disappearance of half of all entry level jobs, reinterpreted the same question a year later as a pure management decision, that companies could do the same with fewer resources, which leads to layoffs, or do more with the same resources. With Amodei the distance this rhetoric has traveled can be measured especially precisely, because as recently as October 2024 he had, in his widely read essay *Machines of Loving Grace*, painted a country of geniuses in the data center whose blessings for health, prosperity and governance he described in glowing colors, while there he expressly left the question of work and meaning unresolved. In only a year and a half the path thus leads from the glowing picture of the future by way of the job loss warning to the soothing management formula, and this swing of the pendulum suggests that the public speech of corporate leaders about the employment consequences of their own technology follows less a stable body of data than the prevailing communicative weather. Amazon founder Jeff Bezos went further still in June 2026 at VivaTech in Paris and held out the prospect of a shortage of workers rather than mass unemployment for the future. A survey by EY-Parthenon among around twelve hundred chief executives in 21 countries underpins this shift with a figure: the share of those who expect substantial AI driven job cuts is said to have fallen from 46 percent in January 2025 to 20 percent in May 2026, though this specific figure could so far be verified only via secondary sources and not on EY's own primary page.²⁷

At precisely this moment of public reassurance, the hard numbers show the opposite. According to the report by Challenger, Gray & Christmas (challengergray.com), in June 2026 artificial intelligence was for the fourth month in a row the most frequently cited reason for job cuts in the United States, with 14,029 AI attributed cuts in that month alone and 101,743 since the beginning of the year, equivalent to around 23 percent of all reported cuts. The technology sector recorded a total of 139,156 job cuts through June 2026, an increase of 83 percent over the same period in the previous year. While the executive boards express themselves more cautiously and more optimistically in public, at the same time companies are justifying more and more of their actual layoffs with artificial intelligence. What is said and what is reported thus diverge, and with that this chapter's thesis of the wrongly named force applies no longer only to individual companies but to the entire industry.²⁷

How little the workforces themselves believe the replacement story that is told over their heads is shown by a survey in which, for once, the employees themselves have their say. The second annual sentiment survey by Lenny's Newsletter, one of the highest reach industry newsletters of the American tech scene, covered more than 5,000 working tech employees in the summer of 2026, predominantly experienced professionals from product management, and as a self selected reader survey it is a mood picture and not a measurement of the labor market, but revealing precisely as a mood picture. Four out of ten respondents are indeed worried about their job, yet the fear of losing one's own job to artificial intelligence ranks at 22 percent near the end of the list of worries, and the connection between the observation that the technology is already taking over parts of one's own work and the fear of layoffs is practically zero. What preoccupies employees instead is the spiral of expectations, that is the worry of having to deliver ever more for the same salary, of getting into a pace that cannot be sustained, and of watching the quality of one's own work decline. The replacement thesis that the corporations offer the capital markets in their communications and that the news agencies perpetuate in their reports on layoffs finds no counterpart precisely where it would have to show up as lived experience, namely in the inside view of those who actually work with the technology. Seen from the inside too, the machine therefore does not act as a job eater but as the metronome of an intensification whose beneficiaries sit elsewhere.²⁸

“The fear of losing one's own job to artificial intelligence stands at the end of employees' list of worries, and the connection between experienced AI takeover and the fear of layoffs is zero.”

The share that the technology can have in these job cuts at all can by now also be bounded from the other side. The economists at Goldman Sachs have extrapolated the use cases currently in operation across the entire American economy and arrive at 2.5 percent of employees whose jobs would be at risk if these applications spread fully, which would lift the unemployment rate half a percentage point above its trend during the transition phase and, by the same calculation, would dissolve again within two years. For those affected that is no all clear, but it is an order of magnitude that bears no relation to the headlines.²⁹

What the extrapolation estimates from above can, since the spring of 2026, be checked from below, against the official statistics. In its Business Trends and Outlook Survey of firms, the American Census Bureau ran a dedicated supplemental module on artificial intelligence with the full sample of 1.2 million establishments between mid November 2025 and early February 2026, recording separately where in the company the technology is used and what changes as a result. The result is sobering for the headline. Of the companies that use artificial intelligence at all, only about one in twenty reports any effect on the number of employees, and these 5 percent are split almost evenly between the two directions, 2.3 percent with an increase against 2.0 percent with a decrease. Where the structure of tasks changes at all, two thirds of companies exclusively expand the tasks of their employees, while pure replacement occurs at only 5 percent. Replacement meanwhile happens elsewhere, because 16 percent of users swap existing software and equipment for AI supported solutions. In this count the technology therefore replaces capital a good three times as often as labor, and that is the opposite of what the public excitement assumes. What

appears in the news as an attack on the workforce is, in the books, first of all an attack on last year's software license. Two caveats belong here: these are self reports by the companies and not a measurement, and the statements about what employees do with the technology come from the company representative and not from the employees themselves.²

How little these measurements reach the public is shown by the annual survey that the polling institute Gallup has conducted together with Bentley University since 2023. In May 2026, 79 percent of adult Americans expected artificial intelligence to reduce the number of jobs in the United States over the coming ten years, six points more than a year earlier and the highest value since the series began. The share of those who ascribe more harm than benefit to the technology overall rose over the same period from 31 to 39 percent, and only 27 percent still trust companies to handle it responsibly. The mood tips most sharply of all among the youngest, whose entry level is the subject of the next chapter: among 18 to 29 year olds the expectation of shrinking employment jumped within one year from 62 to 75 percent, and the share that sees more harm than benefit from 36 to 47 percent, a larger jump than in any other age group. What is remarkable about this movement is its direction, because it contradicts the evidence of this chapter. While executive board rhetoric is softening, while employees in the technology industry barely fear replacement, and while both the extrapolation and the official self reports show the measurable effect to be small, the population is hardening around exactly the story this chapter describes as the AI argument. The job killer thesis no longer needs the layoff statistics, it has become a basic social assumption. An expectation shared this widely is itself an economic quantity, because it shapes career choices, education decisions and political moods long before the machine does what people credit it with.³⁰

“Among 18 to 29 year olds the expectation of shrinking employment jumped within one year from 62 to 75 percent.”

That this expectation runs ahead of the measurement not only in the population but also inside the companies themselves can be read, since August 2026, off a single survey series. In the McKinsey stocktaking of summer 2025, 32 percent of respondents expected the number of employees at their company to fall in the coming year because of artificial intelligence. A year later, in the follow up edition, 14 percent report that this is what actually happened, by the authors' calculation less than half of what was expected, and two thirds state that little or nothing has changed about the size of their workforce. What is remarkable is what follows from this experience, namely nothing. For the year now ahead, 39 percent expect a decline, more than in that previous year's survey whose expectation has just proven too high. The obvious objection, that two different groups are simply answering here, the authors checked themselves and answered in a footnote. For the 552 respondents who took part in both surveys, the finding is the same as for the full sample. The same people therefore missed their prediction and then expect more of it.³

That this expectation behaves differently in small companies than in large ones is shown by the only series in the German-speaking region that has asked it in the same form for three years. For the SME labor market study of the insurer AXA (axa.ch), the research institute Sotomo (sotomo.ch) surveys Swiss companies with five to 250 employees every spring, most recently 336 firms in March 2026. Of the SMEs that say they use artificial intelligence, 12 percent expect to cut staff because of the technology in the coming two years, and this value stands unchanged at 12 percent

in 2024, 2025 and 2026. Just under half expect time savings with no consequences for the number of employees. Only the other side has moved: the share that expects more work because of the technology without adding staff for it has risen from 17 to 22 percent, and the expectation of hiring additional people because of it has halved from 10 to 5 percent. Where there is no share price to move, the expectation of job cuts stays frozen for three years, and the new experience is called extra work. These are self reported answers from a survey that an insurer commissions for its customers, and firms with fewer than five employees, the majority in Switzerland, are missing from it.¹³⁵

This first force is thus entirely real in its consequences for those affected, yet its cause is named wrongly. It is the correction of an overshoot and the cross subsidy of a wave of investment, not the work of a machine pushing people out of the office. Anyone who reads it as proof of the apocalypse confuses the justification with the substance.

This capital market logic has a flip side that is rarely stated openly and that comes, of all places, from the field of corporate valuation itself. The financial economist Aswath Damodaran, who has taught the valuation of young technology companies for decades and who faces the AI boom as a skeptic, laid out in the spring of 2026 that the enormous market sizes assumed on the stock market for artificial intelligence only add up if the technology replaces people and does not merely make their work easier. A tool that makes employees more productive opens up an incomparably smaller market than a technology that makes whole salaries redundant, and for that reason the market size held out in prospect is at its core nothing other than the wage bill of the work it promises to displace. That is a valuation assumption and not an empirical finding, but it explains why the job eater story is cultivated most stubbornly where the most capital hangs on it. Whoever promises the replacement of the human being sells the bigger story, and on the stock market the valuation follows the story. The job killer story is therefore not only the justification for a cost cut, it is at the same time the promise on which a considerable part of today's AI valuations rests, which explains both its tenacity and its vulnerability, because a promise that is not kept gets written off sooner or later.³¹

The same story that dresses up job cuts turns into its exact mirror image on the construction site of artificial intelligence. Where the corporations pass off their cuts as progress, the operators of data centers sell their facilities to regions as a job engine, and both promises hold up worse under examination than their noise suggests. A study by the economists Dany Bahar and Greg Wright, published in May 2026 through the Brookings Institution and deliberately written as an independent academic working paper against the studies financed by the industry itself, linked around 770 American data centers with the employment and wage data of the counties in which they stand and, using the synthetic control method, computed for each of the 93 counties that received their first large data center between 2008 and 2024 a twin without a data center, drawn from around 3,000 unaffected comparison counties. The finding is ambivalent. The facilities do create work, total private employment in an affected county rises by 4 to 5 percent over five to six years, construction gains eleven and the information industry 22 percent, and wages grow by 3 to 4 percent; for a typical county with around 98,000 employees that comes to between 2,000 and 4,000 additional jobs after six years. At the same time the same calculation shows how strongly the industry aligned numbers overstate the effect, because data center counties were already growing

faster than the average before construction, so that a naive comparison of the kind some industry commissioned reports make overestimates the employment effect threefold.³²

The real point for the discussion about public incentives lies in how unevenly the benefit is distributed. Only the large hyperscale facilities, and all the more so entire clusters of four or more data centers, set off the technology effect that is invoked; such a cluster lifts employment in the information industry by 23 percent, while a single facility or a pure colocation center that merely houses other companies' servers leaves barely measurable effects in the information industry. Yet that is precisely where most of the public money flows, because tax breaks and incentives amount to only around 2 percent of the construction investment in counties with hyperscale facilities, but to 62 percent in counties with colocation centers, so that the subsidy weighs heaviest exactly where it sets off the least work. Virginia's tax exemption for data centers alone cost the state more than 1.6 billion dollars in fiscal year 2025, according to its own financial report. For the European discussion, in which the same facilities are courted with the same arguments and the same location incentives, the finding is a warning, even though the numbers come from the United States and the paper is to be read as a working paper that has not yet been peer reviewed. The job engine error is the exact counterpart to the job killer error, and both live off the fact that nobody measures the promised effect against reality.³²

The other side of the discussion deserves the same examination, and it comes out no gentler there. On August 18, 2026, the economists at Landesbank Hessen-Thüringen assembled the reasons why artificial intelligence is not to be expected to produce high unemployment for the economy as a whole, and named four mechanisms that are regularly missing from the doom scenario: productivity gains make human work more valuable instead of making it redundant, automation frees up capacity for other activities, innovations bring forth occupations that nobody could name before, and higher incomes generate demand that in turn creates employment. These mechanisms are well founded, the analysis expressly draws for them on a 2015 essay by David Autor, and the official American business survey above supports them involuntarily, because there capital is replaced a good three times as often as labor. Only the example with which analyses of this kind have been supporting their all clear for years has itself by now become a case that has to be checked against the time series. The example cited is American bank tellers, whose number had risen by 50,000 to around 550,000 between 1980 and 2010 despite the spread of ATMs. According to the occupational statistics of the American Bureau of Labor Statistics it stood at 347,400 in 2024, and the same agency expects a further decline of 13 percent by 2034. The mechanism remains correct, its piece of evidence does not. Anyone who supports the all clear with a snapshot whose time series has shown the opposite for a decade and a half makes the same mistake as the headline he wants to refute, namely confusing a story with a measurement.³³

That this argument can be invoked not only for job cuts but also for one's own added workload has been demonstrated since August 2026 by a German agency. Andrea Nahles, chair of the executive board of the Federal Employment Agency, describes in an interview with Deutschlandfunk on August 15, 2026 an increase in appeal proceedings against job center decisions and traces it to the use of artificial intelligence by those affected. Where half a page used to be enough, ten pages are now common, which in her words creates an insane amount of extra work for the staff, especially since the longer letters sometimes contain hallucinations, that is, invented citations and court rulings. For a single federal state the development can be quantified, because the Regional

Directorate North reports a total of 10,348 appeals for Schleswig-Holstein in the first half of 2026, which extrapolated is around a quarter more than in the previous year. How many of them actually came about with machine assistance the agency expressly cannot say, because that is recorded nowhere.³⁴

It is exactly at this point that a look into the official statistics pays off, and it reorders the case. According to the agency's annual analysis from January 2026, 501,667 appeals were already filed nationwide in 2025, 78,310 more than the year before, and the appeal rate of the joint institutions rose from 1.8 to 2.1 percent across around 21.2 million decisions, with the agency naming as the most frequent points of dispute the calculation of housing costs, the withdrawal and reimbursement of benefits, and the crediting of income and assets. The increase had therefore begun long before anyone attributed it to the technology. The number with which the case is told publicly also measures something other than what it appears to measure. Of the 476,728 appeals settled in 2025, 147,213 led to a changed decision, from which the reporting makes the sentence that more than one in four appeals succeeds. But the job center had decided in legal error in only 42,303 cases, while almost half of the changes came simply from documents being submitted later in the proceedings or outstanding duties to cooperate being fulfilled.³⁵

What remains of the case nevertheless carries weight, because it contradicts the common job killer story. If the agency's observation is correct, then in public administration the technology first appears not as work saved but as additional work, set off from the other side of the desk, and it hits precisely that sector which has the oldest workforce in the DACH region. Nahles' first idea, letting artificial intelligence talk to artificial intelligence, she rejects in the same interview as a permanent solution, and for good reason, because it describes a cycle in which both sides automate their effort while the procedure becomes neither faster nor more correct.³⁴

Looking ahead

If the mechanics of this chapter are right, then the AI argument has an expiration date. Capital markets reward a story only for as long as it cannot be checked, and exactly this checkability grows with every quarter in which the missing link between cuts and success established by Gartner holds.¹⁸ It is foreseeable that investors will begin to ask the question this study asks, namely which concrete task an AI system has actually taken over when job cuts are justified with the technology, and companies that have no answer to that will then pay more for this argument than it earns them today.

The real opportunity lies on the other side of the same calculation. The billions that job cuts free up today do in fact exist, and the companies that direct them not into the story but into the patient rebuild of their processes are buying exactly that capability which becomes visible in the further course of this report as the truly scarce resource. In a few years today's balance sheets will show who understood the wave of layoffs as a message to the stock market and who understood it as an investment in a new organization.

The career gap is measurable, but its cause is not so easy to pin down

The second force is quieter than the first and considerably more serious, because it does not hit the experienced, who are reorganizing themselves, but the young, who have not even started. It shows up in the harder entry into working life, and unlike the wave of layoffs it cannot be dismissed as a mere justification. Here there is a real and measurable effect. Only its cause is anything but settled, and it is precisely this lack of clarity that is analytically the most interesting part of the whole discussion.

Before the individual studies get their say, it is worth looking at the measure they all work with, because in the public discussion it is misunderstood almost without exception. When the talk is of occupations strongly exposed to artificial intelligence, this exposure captures how many of an occupation's typical tasks overlap with what the technology can do today, and precisely not how likely it is that these activities will be replaced. The OECD sharpened this distinction once more in its most recent paper on skills in the age of AI, and in doing so drew a picture that looks contradictory at first sight. The most exposed of all are the highly qualified white collar occupations, the IT specialists, the business professionals, the managers and the engineers, that is, the activities that work most with language, analysis and knowledge. Yet the same occupations carry the lowest risk of actually being automated, because they rest on non-routine judgment, on social coordination and on experience that has so far escaped the machine, while conversely the highest automation risk lies with the low and medium skilled activities with routine processes, which sit far down the exposure ranking. Exposure therefore describes where the technology can change and complement work, and not where it displaces people, and anyone who equates the two is already committing that confusion of observation and cause which this chapter is about. An honest reading includes the fact that the underlying assessment of automation risk cannot yet capture the full effect of the widely available generative tools, because its survey period predates them, so that the line between mere exposure and real replaceability remains in motion.³⁶ Since August 27, 2026 there has been official evidence for this. For the first time, the Bureau of Labor Statistics has attached its own exposure ranking to its employment projection for the years 2025 to 2035, and has furnished it with an unusually strict caveat: the tier is not an employment prediction, not a wage prediction and not an estimate of whether people will be replaced, and it does not even distinguish between automation and complementation; the ranking measures no employment effect, it merely supplies additional information for choosing an occupation. How seriously this caveat is meant can be read off the projection itself. Nursing sits in the second highest exposure tier and gains around 194,700 positions by 2035, software development sits in the highest tier and adds a good 10 percent, while in that same highest tier general office clerks lose 156,200 positions. Whoever ranks occupations by their exposure therefore gets a different order than whoever ranks them by expected demand.³⁷

The widely noted starting point is a study by the Stanford Digital Economy Lab that rests on actual payroll data of American employees and that the authors have updated twice since; the


authoritative version is the one of August 12, 2026, with data through June 2026. It finds that employment of 22 to 25 year olds in those occupations that are particularly exposed to artificial intelligence lies 19 percent below the level it would have reached had it followed the growth of their less exposed peers. In absolute numbers, employment in this age group fell by around 11 percent in the two most exposed fifths between the end of 2022 and June 2026, while it rose by around 10 percent in the three least exposed fifths. Older and more experienced employees show no comparable shortfall in the same occupations; the 35 to 49 year olds in the two top exposure fifths even gained around 10 percent. At first sight the technology here seems to be doing exactly what it is said to do, namely closing off the entry level and rewarding experience. What matters is that this is a relative effect between more and less exposed occupations and not a measured loss of jobs in absolute numbers.⁴

The real gain of the new version, however, lies not in the number but in the question of how this gap comes about. A shrinking headcount can have two causes, fewer new hires or more separations, and the two mean something entirely different. The authors separate this cleanly for the first time and find that separation rates among young employees in the exposed occupations fell at least as much as in the weakly exposed ones, that is, in exactly the opposite direction to what displacement would mean. The entire gap arises on the hiring side. The finding is therefore not the picture of a machine pushing people out of their jobs, but of an entry that is narrowing, and this distinction says more than any percentage figure. At the same time, the authors set out the weaknesses of their own measurement more openly than in any earlier version. The gap shrinks from 18 to 9 percentage points as soon as the education level of the occupations is netted out, and they themselves leave open whether education is a competing explanation here or precisely the channel through which the technology works. And it turns out considerably smaller in the official American household survey than in the payroll data, at minus 0.022 against minus 0.132 for the same window, while both sources almost agree in the information industry, financial services and business services, at minus 0.213 against minus 0.231. The direction of the finding is thus supported more broadly than ever before, but by the authors' own judgment its magnitude is tied to their data source and cannot be scaled up to the economy as a whole.⁴

Since the summer of 2026, however, this finding can no longer be dismissed as a snapshot. In June 2026 the same lab, together with the research unit of the payroll provider ADP (adp.com), launched a monthly updated dashboard that carries the measurement forward as a running series, and the first research note on it shows a gap that keeps widening instead of closing. In the underlying sample of 25,000 firms covering more than 700 occupations, employment of 22 to 25 year olds in the most exposed occupations has been shrinking by 3.8 percent per year since the launch of ChatGPT, while the least exposed occupations in the same age group are growing by 2 percent per year, and according to Erik Brynjolfsson, the lab's director, the decline has been deepening by around half a percentage point per month since the first publication. Muted signs of the same pattern now reach into the age group up to 34, which in the original study had still been unremarkable. The finding itself thus stands more stable and more current than ever before, and its interpretation remains untouched by that, because the updated measurement too reports a relative gap between occupational groups and no quantified job loss.³⁸

The gap at the entry level widens in the exposed occupations.

Employment June 2026 against Nov. 2022, in percent



• Gap of the youngest to the expected level: **19 percent.**

Source: Brynjolfsson, Chandar, Chen, Canaries in the Coal Mine, Stanford, version of August 12, 2026, ADP data, US only.
Note: Difference between occupational groups, not a measured job loss.

FIGURE 3

Exactly at this point the scientific controversy begins, and it is manned with weighty arguments on both sides. A study from the Federal Reserve by Jessica Liu and Douglas Webber ("AI Adoption and Firms' Job-Posting Behavior", March 2026), which evaluates more than a million companies, yields no indication that the adoption of artificial intelligence depresses employment, and with estimates so tightly bounded that even a small effect would have been noticed; what is measured there is job posting behavior and not wages themselves.²⁰ The study by Anders Humlum and Emilie Vestergaard ("Still Waters, Rapid Currents", 2026), which rests on Danish register data covering 25,000 employees from 7,000 workplaces, likewise finds no effect on income or working hours; it does confirm the collapse at the entry level, but shows that it does not originate with those firms that actually adopt the technology.³⁹ The collapse is therefore demonstrable, yet attributing it to artificial intelligence stands on shaky ground.

A special place in this dispute belongs to a research note that Anthropic (anthropic.com) itself put out in March 2026, written by the economists Maxim Massenkoff and Peter McCrory. It develops a new measure that interlocks the theoretical feasibility of an occupational task with the actually observed occupational use of its own models, weighted by degree of automation, and then tests the employment consequences with the official American household survey, that is, with a third data source alongside the payroll records of the Stanford series and the job postings of the study by the Federal Reserve. The result turns out remarkably sober. Since the end of 2022 there is no systematic rise in unemployment in the most exposed occupations, no matter where the threshold for exposure is drawn and in the unemployment insurance data as well, and the unemployment rate of the youngest stays just as flat. Only on the hiring side does the familiar pattern repeat itself, because the rate at which 22 to 25 year olds take up a new position in a strongly exposed

occupation lies 14 percent below the 2022 level after the launch of ChatGPT, by the authors' own judgment just barely statistically significant, while among older workers no such decline can be observed. The direction of the Stanford finding is thus replicated once more and in a further data source, but distinctly more weakly, and the weight of this source lies precisely in its origin. A vendor whose valuation hangs on the story of its technology's economic force would have every reason to report dramatic employment effects, and instead publishes a null finding on unemployment together with a finding on the entry level that is secured by the thinnest of margins. As a self-measurement on its own platform, the note nevertheless remains driven by interest and cannot be recomputed by outsiders, and the authors themselves record that the young who are not hired may also stay in existing positions, take other paths or return to university.⁴⁰ Exactly this weakness is offset by a finding that comes from a quite different corner. Sarah Eckhardt and Nathan Goldschlag, working for the Economic Innovation Group, divided American workers into five equally sized groups by their degree of exposure and tracked the development of their unemployment rate between 2022 and the beginning of 2025. In the fifth of the most exposed employees the rate rose by 0.30 percentage points, in the fifth of the least exposed, by contrast, by 0.94. Anyone looking for the traces of a machine pushing people out of employment finds them weakest of all precisely where they ought to be clearest. This evaluation too comes not from a peer reviewed journal but from a Washington think tank, yet it has no vendor behind it and it measures unemployment where the Anthropic note remains dependent on its own platform.⁴¹

An even more recent study shifts this dispute onto a third level of observation and thereby makes it clearer rather than settling it. The research team of the financial services provider Ramp linked the actually observed spending behavior on artificial intelligence for a good 21,000 American firms with the employee data of a labor market data vendor and found that after adoption the companies expanded their headcount by a good 10 percent over the following two years, but that this increase was borne almost entirely by the most intensive adopters, while the hesitant adopters showed no appreciable change. What is remarkable is that the entry level of all things grew fastest, because at the most heavily investing companies the number of entry level employees rose by 12 percent over two years. At first sight this finding cuts across the Stanford result, according to which the entry level is closing off for the young, and yet the two do not contradict each other once one recognizes that they ask different questions. Stanford examines occupation and age and answers in which activities access is narrowing for the young, while Ramp examines the firm and the intensity of its adoption and answers which companies grow. A company can hire more people in total because it scales with the technology and becomes more productive, and in the same move devalue the entry level within the particularly exposed occupations. The Stanford researchers even record this separation themselves, in that they recompute their collapse once more without the information sector as a check and find it almost unchanged, so that it does not hang on the technology occupations, while the gains at Ramp conversely cluster precisely in the information sector and at the most investment intensive firms. This finding too calls for the customary caution, because the source is a commercial vendor with its own interest in rising AI spending, because it measures a correlation and not a causal build up of jobs, and because its adopters are from the outset the larger and faster growing companies. What it contributes to the discussion is therefore not an all clear, but the insight that the second force is not to be read as an either or, but as a coexistence of two axes on which the young can lose in relative terms while the adopting firms grow in absolute terms.²¹

A third decomposition of the same finding was put forward in August 2026 by a quantitative macro fund, and it starts neither from the occupation nor from the firm but from the individual work task. Benjamin Verschuere and Angus Cameron of Liminal Capital (liminal-capital.com) had language models sort all 17,536 task descriptions of the American occupational database O*NET into three classes, work that the machine completes on its own from start to finish, work for which a human has to be present or has to sign off, and physical work that it does not touch, and then mapped the result back onto occupations through the official employment and wage statistics. On that basis, 32 percent of American employment falls into the first class, 15 percent into the second and 53 percent into the third, and since 2021 work in the first class has grown 3.3 percentage points a year more slowly than the untouched work, while work in the second class has grown faster by the same margin. The figure that matters for this chapter stands at the beginning of the working life: the number of 22 to 25 year olds taking up a job in the most replaceable occupations is roughly a third below its 2021 level, by the authors' calculation around 227,000 entry level positions a year, without layoffs in the same occupations having increased. The New York Times put this finding on the front page of its business section on September 16, 2026, under the headline "The Quiet Way A.I. Is Hitting the Work Force". The value of this work, however, lies not in the headcount, which merely confirms the familiar pattern once more, but in the fact that it sets replaceable and completable shares against each other within the same occupation instead of booking an occupation as exposed as a whole. That is exactly why the averages of the large exposure indices show so little, because in them two opposing movements cancel each other out that in reality hit different people. An honest reading includes the fact that the calculation here comes not from a research institute but from a fund that lives on recognizing such turning points early, that the work is not peer reviewed and that the classification of tasks comes from the very technology whose effect is to be measured; the authors themselves describe the results of their model as a scenario and not as a forecast.¹³⁶

That the collapse is being attributed to the technology with growing matter of factness despite this uncertain body of evidence is demonstrated by the Corporate Recruiters Survey 2026, a worldwide survey of hiring managers that the Graduate Management Admission Council (GMAC), the association of leading business schools that administers the GMAT admission test, published in early summer 2026. More than 600 hiring managers from 39 countries were surveyed, a good half of whom recruit for companies on the global Fortune 500 lists, and around a third of them stated that they are replacing entry level positions with artificial intelligence, led by the technology industry, where 40 percent made this statement, followed by manufacturing. At first sight the cause finally seems to be named here, because the employers themselves cite the machine as the reason, yet precisely this form of evidence calls for the greatest restraint. It is a self-report about what hiring managers believe they are doing, and not a measurement of what the payrolls show, and it comes moreover from an association that has its own interest in the value of a business education. The figure therefore supports the direction of the Stanford finding, but it is at the same time the purest example of that confusion of observation and cause which this chapter is about, because a real hiring freeze at the entry level is furnished with the justification that promises the most attention in the current climate.⁴²

An obvious and uncomfortable alternative explanation is the sharpest turn in interest rates in four decades, which made every new hire at the entry level more expensive overnight. It explains why

the entry level in particular suffers, without any machine having to be involved, because entry level positions are the first that a company stops filling when money becomes expensive. There is a precedent that demonstrates this pattern strikingly. In the phase after the bursting of the dot-com bubble, a widespread story with no empirical basis kept young people from studying computer science, until the market recovered a few years later and demand returned. The pattern is repeating itself with astonishing precision. A steep rate increase is followed by a hiring freeze at the entry level, a convenient technological scapegoat supplies the explanation, and as soon as rates fall, demand returns.

Since January 2026 this interest rate explanation has also been available in worked out form, and its origin makes it doubly remarkable, because it comes from Zanna Iscenko and Fabien Curto Millet, two economists from the staff of the chief economist of Google (google.com), published through the Washington think tank Economic Innovation Group. Their evaluation of more than 238 million American online job postings shows that postings in the most exposed occupations already reached their peak in the spring of 2022 and fell steeply thereafter, more than half a year before ChatGPT appeared and in exactly the month in which the Federal Reserve began its sharpest series of rate increases in four decades. The employment decline from the end of 2022 would accordingly be only the delayed echo of a hiring freeze that came before it, because normal turnover is no longer offset by new hires. To this is added a mechanical argument that bears directly on the construction of the Stanford measurement. An age group as narrow as that of the 22 to 25 year olds is, unlike the older cohorts, not refilled from below by cohorts moving up, but lives almost entirely on new hires, and in a worked example by the authors a single year of hiring freeze without a single layoff is enough to shrink it by a quarter, while all other age groups remain unchanged. Finally, the exposed occupations cluster in the most interest rate sensitive industries of all, because around 38 percent of the employees in the top exposure fifth work in the information industry, in finance and in business services, against less than 2 percent in the bottom fifth. That of all people the chief economist of a technology group is the one presenting the reassuring reading has to be stated openly and counsels caution, but it does not devalue the data, and so the interest rate thesis has since stood in the room no longer as a suspicion, but as a worked out counter thesis.⁴³

The Stanford researchers replied to this in February 2026 with an analysis of their own, and their answer shows how much this dispute sharpens the discussion instead of paralyzing it. AI exposed occupations, runs their first objection, are on average precisely not the most interest rate sensitive ones, the two measures are even negatively related, and the shortfall of the exposed occupations shows up regardless of how interest rate sensitive an occupation is. With the second objection the researchers turn the timing argument against the interest rate thesis, because in the calculation with the most comprehensive controls, which fully nets out firm specific shocks, the collapse among the youngest only becomes statistically detectable in 2024, that is, exactly in the period in which companies actually built the technology into their processes, and thus too late for the turn in interest rates and too early for an all clear. This second objection, however, is to be read with a caveat, because that very control calculation is one the authors report in their version of August 2026 only as consistent in direction, with a precision that fluctuates across the subsamples. At the same time the authors concede that artificial intelligence is unlikely to be the sole cause of the development. The dispute is therefore not decided, but it has become more precise, and precisely

this sharpening is the gain, because the question is no longer whether the career gap exists, which the other side does not dispute either, but how much of it belongs to the machine and how much to expensive money.⁴

Since late spring 2026 a second alternative explanation has stood alongside the turn in interest rates, and it is the analytically sharpest challenge that has so far arisen to the AI attribution, because it does not merely voice the suspicion but tests it. A working paper from the London School of Economics with the telling title *The Broken Ladder*, written by Peter John Lambert and Yannick Schindler, starts from a simple observation: the occupations most exposed to generative AI are for long stretches the same ones that can most easily be shifted to remote work, so that any study that looks only at AI exposure involuntarily measures the second great post-pandemic shock as well. On the basis of around 240 million new hires and a good 400 million job postings from the United States, Great Britain, Canada and Australia for the years 2017 to 2025, the authors first show the familiar finding, because considered on its own a markedly raised AI exposure goes together, up to 2025, with a junior share of new hires that is around five percentage points lower. Only the same holds for remote work exposure, and as soon as both quantities are estimated jointly, the remote work effect persists, while the AI coefficient shrinks sharply and is often no longer statistically distinguishable from zero. That in these data the collapse at the entry level already sets in during 2022, and thus before the mass diffusion of the generative tools, fits this picture. The authors do not claim that the technology plays no role, but they do claim that a considerable part of what the discussion attributes to the machine is in truth a consequence of distributed work, which has made the training of beginners expensive, and this paper too, as a working paper that has not yet been peer reviewed, is to be read with the appropriate caution.⁴⁴

Almost at the same time, a post on the research platform of the Federal Reserve Bank of New York, written by Natalia Emanuel, Emma Harrington and Amanda Pallais, supported the same reading from the unemployment side. The unemployment rate of young American college graduates rose by a fifth between the years before and after the pandemic, while among experienced graduates it even declined slightly, and this widened age gap sits entirely in those occupations that can be done well remotely, while in in-person occupations it disappeared again long ago after a brief spike in 2020. By the authors' back-of-the-envelope calculation, the shift to working from home thus explains 64 percent of the entire increase, and they describe the mechanics behind it using the internal data of a Fortune 500 corporation: those who sit next to their colleagues receive measurably more feedback and mentoring, the loss hits the youngest hardest, and the same firm permanently filled positions in distributed teams with experienced people, even after the offices had long been open again. Here too caution belongs to the picture, because this is a blog post with a calculation expressly labeled as an estimate and with mechanics from a single software firm, and the authors themselves note that the technology may have a stronger effect in the future. For the question of this chapter the point is considerable nonetheless. The dispute over the cause of the collapse at the entry level thus has three candidates, the machine, expensive money and distance, which do not exclude one another, and the most honest reading is that the most popular of the three explanations is currently the one with the weakest empirical support. Yet the explanations converge in their end point in a troubling way, because whether the entry level loses its learning function to the machine, which takes over the routine preparatory work, or to the distance, which cuts off mentoring, it thins out in either case, and the question of where the experienced people of

the day after tomorrow are supposed to come from arises regardless of which of the three candidates ultimately proves to be the main cause.⁴⁵

“The unemployment rate of young college graduates rose by a fifth between the years before and after the pandemic, among the experienced it fell, and this age gap is found entirely in those occupations that can be done well remotely.”

Since the spring of 2026 there has for the first time been a study on this dispute that tests both candidates against each other instead of dealing with them one after the other, and it comes from Europe. Magnus Lodefalk, Lydia Löthman, Michael Koch and Erik Engberg of the AI-Econ Lab at Örebro University use a coincidence of the Swedish interest rate cycle as a measuring instrument. The Riksbank began its rate increases in April 2022, seven months before ChatGPT appeared, and this gap makes it possible to separate the two explanations in time, which is not possible in the American data because the two ran almost simultaneously. The result is split in two. In the 4.6 million job postings of the public employment service the decline clearly goes back to the turn in interest rates, the effect of the ChatGPT launch is statistically indistinguishable from zero, and it clusters of all places in those occupations that precisely cannot be moved to working from home. For employment itself a different picture emerges. On the basis of the complete population register, that is not a sample but all employment relationships with monthly employer reporting, an age gap opens up within the same employer. In the most exposed occupations, employment of 22 to 25 year olds in the first half of 2025 is 5.5 percent below that in less exposed occupations, for 26 to 30 year olds it is just under 5 percent, and for those over 50 there is an increase of a good one percent. Among young women the decline is about twice as large as among young men, and a third of this gap disappears once their occupational distribution is statistically adjusted to that of the men, while the rest sits within the same occupations. This is the first evidence outside the United States, and on register data rather than on surveys or postings, for what this chapter claims: overall demand is not collapsing, but who gets hired is shifting. The authors themselves urge caution. The paper has not yet been peer reviewed, the 5.5 percent is the end point of a curve that diverges over time and not the average effect, which is considerably smaller, and from Finland there is a null finding using the same research design. That is exactly where things stand: the effect is there, it is now documented in Europe as well, and it is smaller and slower than the headlines present it.⁴⁶

How the technology affects the entry level can by now be measured. The earlier value of an entry level hire lay in being the cheap labor for the routine preparatory work, for the research, the first draft, the simple code. It is precisely this preparatory work that the digital team member now takes over. That this channel is actually at work has been shown by the Stanford researchers, who used Anthropic's usage index to break usage per occupation down into an automation share and an augmentation share, and only the first is connected with the employment trends. The larger the share of tasks that are handed over entirely to the machine, the weaker employment of the youngest in an occupation develops, while the share of genuine collaboration between humans and the AI colleague shows no discernible connection.³⁸ What retains its value is the judgment that only experience confers. In this way the technology upgrades experience and devalues the entry level without displacing the experienced.

“The technology upgrades experience and devalues the entry level without displacing the experienced.”

A recent Austrian study backs this pattern with a concrete number. An analysis of around 25,000 open positions published in June 2026 counts, among the more than 1,000 postings classified as data or AI positions, only about 6 percent that are aimed at people at the entry level, while almost 80 percent address experienced profiles and the ratio of senior to junior is more than six to one. This number is valuable as evidence, but it demands the same care that this study applies everywhere. It comes from a commercial source, it measures job postings and not actual hiring, and it captures a single reference date, so that it can show that the entry level share is low today, but not that it has fallen. When the accompanying coverage nonetheless attributes this low share to artificial intelligence, it commits exactly that confusion of observation and cause which this study criticizes, and thereby unintentionally becomes an example of the described pattern itself.⁴⁷

A German finding close to official statistics at the same time sets a limit to the story of an entry level that is uniformly without prospects, and it also demonstrates how much a single aggregate number conceals. According to the analysis by Timon Hellwagner and Enzo Weber at the Institute for Employment Research, based on the microcensus and the statistics of the Federal Employment Agency, the labor force participation of 20 to 24 year olds in Germany stood at 76.8 percent in 2025 and thus at its highest level in decades, an increase of 7.1 points since 2015, while 25 to 64 year olds gained only 3.1 points over the same period. Anyone who reads this as an all clear for the entry level, however, is measuring the wrong quantity, and that can be shown from the same source. For one thing, according to an earlier analysis by the same institute, the increase between 2015 and 2023 comes predominantly from students, whose participation rate rose by 19.3 points to 56 percent, while among non-students it rose by only 1.6 points to 85.9 percent. To a considerable extent what is measured here is therefore the side job alongside studying and not entry into a profession. For another, the 2026 analysis shows where this labor force participation sits. Young people are overrepresented by 4.4 percentage points in food and hospitality occupations and by 3.4 percentage points in retail occupations, while they are underrepresented by 3.0 percentage points in business related service occupations and by 3.8 percentage points in company management and organization, that is precisely in those office occupations over which the dispute about AI exposure is conducted. Both caveats belong with it: the authors themselves reduce the lead in hospitality to 2.7 percentage points once part-time students are taken out, and the fact that 20 year olds rarely sit in company management is not an effect of the technology but a matter of course in any career. What remains nonetheless fits exactly into the picture that the Swedish register data draw. Labor force participation as a whole is high and continues to rise, and yet within the exposed occupations who gets hired is shifting. The one does not refute the other, they are two different measurements, and only the second answers the question of this chapter.⁴⁸

A third German measurement demonstrates that the entry level is also narrowing where the technology is certainly not the cause. Andrea Nahles, chair of the executive board of the Federal Employment Agency, reported in mid August 2026 on 433,000 apprenticeship positions registered for the current year, 7 percent fewer than in the previous year, of which 162,000 were still unfilled at that point, and on the fact that only 55 percent of firms train apprentices at all any more. As the reason she names not artificial intelligence but companies' uncertainty about order books and

economic prospects. For the question of this chapter that counts twice. In the German speaking countries the entry level runs to a considerable extent through the dual apprenticeship system and not through the academic office job on which the American discussion hangs, and it is precisely this path that is thinning out for cyclical reasons. Added to this is what, by the same agency's figures, hangs on it, because anyone without a completed vocational qualification carries an unemployment risk of 21 percent over their working life, compared with 3 percent for those with a qualification. Anyone who tells the career gap solely as an effect of the machine overlooks the fact that the most reliable protection against its consequences is already narrowing for other reasons.⁵

How sharply entry via vocational training differs from entry via the academic office job can, since September 2026, be read off a time series for Germany for the first time. The Hiring Lab, the research arm of the job board Indeed, has sorted the job postings on its own platform since 2019 by whether their title signals an entry-level position, an internship or an apprenticeship, and the three series have diverged since 2022. Postings for entry-level positions peaked in July 2022 at up to 75 percent above the 2019 average, have since fallen by more than two thirds and in July 2026 stood 46 percent below the 2019 level, while the market as a whole, after a decline of 38 percent, was still 8 percent above 2019. In 26 of 28 occupational groups the supply of entry-level positions fell, in 23 of them more steeply than the overall market, by 75 percent in marketing, 70 percent in software development and 77 percent in IT applications; it grew only in nursing, medicine and pharmacy and in gastronomy and hospitality, that is, in occupations with staff shortages that do not require a university degree. Postings for vocational training including dual study programs, by contrast, stood on average 50 percent above the 2019 level in the season from August 2025 to July 2026, despite slight declines since 2023, and in 20 of 23 occupational groups the supply of apprenticeships developed better than entry-level positions; in software development, 70 percent fewer entry-level positions stood against 27 percent more apprenticeships for IT specialists. Two caveats belong here, and the author of the analysis, Virginia Sondergeld, draws them herself. Apprenticeship occupations lie to a greater extent in fields the technology barely reaches, since only in 15 percent of apprenticeship postings can more than 60 percent of the required skills be taken over or supported by generative AI according to the company's own measure, against 42 percent of entry-level postings, and for companies in uncertain times it can simply be the cheaper option to train young staff themselves rather than hire university graduates. What the analysis says about causes fits the picture of the three candidates: The gap between entry-level positions in heavily and in lightly exposed occupational groups opened in the spring of 2022, half a year before ChatGPT, and the Hiring Lab states that an AI effect cannot so far be demonstrated in the data. What is measured are postings on a single platform, filtered by words in the job title, not hires and not employment.¹⁴⁰

How these data land in public view became apparent even before the analysis was published. Die Zeit ran them on September 18, 2026 under a headline about where entry-level jobs are disappearing because of AI, and in the text let Melanie Arntz, Vice Director of the Institute for Employment Research (IAB), explain the opposite: In an economic crisis it is usual for entry-level positions to disappear, because companies are reluctant to lay people off, rely instead on age-related departures and hire fewer new staff, which means less effort and keeps knowledge inside the company. That is, in one sentence, the mechanism this chapter has measured in American and Swedish data: The gap arises on the hiring side and not through separations, and it hits the

youngest because their age group consists almost entirely of new hires. That this restraint keeps knowledge inside the company is at the same time the reason why the experienced retain their value, and it leads to the question this chapter asks at the end, namely where the experienced of the day after tomorrow are supposed to come from.¹⁴¹

What these German measurements show for positions and apprenticeships has had a counterpart in wages since August 2026. In its June 2026 business survey, the Munich based ifo Institute asked more than 3,000 German companies that already use artificial intelligence how the technology will affect their employees' wages over the coming five years, and the answers diverge markedly by professional experience and qualification. For workers with less than five years of professional experience about every second company expects falling wages, for the more experienced around four in ten. The most revealing number in the survey is the one that contradicts the common expectation. For entrants with a university degree the share, at 50.8 percent, is even slightly above the 48.3 percent for entrants without a degree, and the prospect of a rising wage arises only with professional experience and then above all for the graduates, for whom 26 percent of companies consider a positive wage effect possible, compared with 16.3 percent for graduates at the start of their careers and 5.9 percent for entrants without a degree. The formal qualification therefore does not protect the entry level, it pays off only together with experience. This is exactly the shift that this chapter has already measured in postings, in stocks and in hiring, here for the first time from German firms and not in positions but in wages. Honesty requires noting that expectations were surveyed and not wages actually paid, that only companies already using the technology answered, and that the institute publishes neither the wording of the question nor the answer categories, which leaves open whether the respondents were thinking of wages actually falling or of slower growth.⁴⁹

Since the spring of 2026 the Austrian snapshot from the job postings can be checked against a second, independently collected measurement that counts not postings but the stock. A research team at the Vienna University of Economics and Business around Christian Schumacher and Can Tihanyi, together with Andreas Schumacher, the chief innovation officer of the plant engineering firm VTU, measured the Austrian AI workforce from the standardized occupational profiles of Revelio Labs, that is essentially from the data of professional networks such as LinkedIn and XING, which cover a good fifth of official employment: 2.86 million position records from eight years.⁵⁰ By this account Austria has around 4,600 specialists visible in the panel whose occupation consists at its core of developing artificial intelligence, operating it or embedding it into business processes, and of them only 6 percent are at the entry level, compared with 23 percent in the rest of the employment in the same panel. The stock thus confirms the postings, and it says more than they do, because it shows how this workforce came about: not through the entry level but through lateral moves by experienced specialists from adjacent technical and analytical occupations. There is a second finding on top of that which sharpens the career gap. The few who do enter are the most likely to leave again, a good fifth of the entrants continue their career abroad, at the top of the hierarchy it is just under one in ten. The authors themselves read the 6 percent in two ways, as a characteristic of an experience intensive activity and as a warning that a workforce without new talent ages. For this chapter the point lies elsewhere: of all places, it is in the occupations that the technology itself brings forth that the entry level is narrowest. The report has not been peer reviewed and rests on a panel of professional network profiles, its absolute numbers are therefore

panel sizes and not official stocks; the pattern, however, matches what postings, American data and Swedish data show alike.

An unusual piece of evidence for the same pattern comes from a June 2026 study by Harvard Business School and INSEAD that does not examine existing workforces but newly founded firms themselves. Hyunjin Kim and Rembrand Koning linked the employee data of around 2,200 American start-ups from the cohorts of the accelerator Y Combinator (ycombinator.com), as well as of a broader sample of more than 40,000 venture financed firms, with a database of employee profiles and found that AI native companies are around 25 percent smaller than comparable non-AI firms of the same founding cohort and the same industry and, on average, organized half a hierarchy level flatter. At the firm level the share of entrants at the AI native companies is around 4 to 6 percentage points lower and the share of the experienced correspondingly higher, a pattern that corresponds exactly to the shift described earlier within the same occupations, observed here, however, not at existing firms that are rebuilding themselves, but at firms that are designed from their founding onward with less entry level staff. The authors themselves expressly urge caution about translating their firm level finding directly into an economy wide employment finding, because artificial intelligence at the same time makes far more company formations possible than before, the number of first time AI financings in the broader sample was around eight times as high in 2024 as in 2020, so that a firm that is smaller on average does not necessarily mean less total employment in the segment. For the second force under discussion here it nonetheless remains notable that the devaluation of the entry level shows up at two independent levels of data, once in the wage data of existing occupations, once in the organizational structure of newly founded firms.⁵¹

“AI native companies are around 25 percent smaller and organized more flatly than comparable firms of the same founding cohort and the same industry.”

That this pattern does not remain confined to existing occupations and newly founded firms is suggested by an observation from the summer of 2026 that comes from practice and can count only as an observation from practice. The American content marketing specialist Kaleigh Moore counted around 50 advertised full-time positions for visibility in AI search systems within two weeks, a job profile that did not exist 18 months ago and that is only now looking for a name under abbreviations such as AEO or GEO, at companies from Stripe (stripe.com) through Pfizer (pfizer.com) to eBay (ebay.com) and across the industries. What is revealing about this collection is less the number, which comes from the private count of a commercially interested observer and measures no market size, than the level at which the positions are located. The advertised positions are almost without exception roles for directors, vice presidents and senior managers with their own budget and their own reporting lines and with salary ranges that reach more than \$200,000 at the top, while entry level and junior positions are practically absent from the collection. The first marketing job profile that the AI wave itself has brought forth is therefore being staffed from the top down, and where new work arises the entry level is apparently not created in the first place. Job postings prove no employment effect, but they do prove organizational decisions, and this decision turns out here the same way as at the AI native new firms: the new field begins with experience.⁵²

“The first marketing job profile that the AI wave has brought forth is staffed from the top down, entry level positions are not even provided for.”

How far this pattern has by now reached established market observation is shown by a prediction from the market research firm Gartner dated July 2026, which is doubly revealing for the question of this chapter because it unites both sides of the calculation in itself. On the one hand Gartner expects that in 2029, 60 percent of organizations will move their software development across the board to considerably smaller teams, compared with 15 percent in 2026, to units of today typically four to five and in future in part only two to three people, in which the classic role boundaries blur and every member, alongside their own specialty, also carries the oversight of AI agents. On the other hand Gartner precisely does not connect this expectation with shrinking workforces, because the responsible analyst Aliyah Camacho expects the technology to make demand for software developers rise overall and not fall, because the growing need for software will exceed the efficiency gains, and she insists that the small teams are not a cost cutting tactic but a reordering that aligns itself with the respective strengths of human and machine. The real point for this chapter lies in the warning that Gartner attaches to the same release. Those who rely on the technology in order to cut junior positions, it says, impede the transfer of knowledge, thin out their own supply of new talent and in the end force themselves to recruit in the most expensive and most contested segment of the labor market, and by 2028, according to the dated prediction, organizations that proceed this way will have hollowed out their own base of new talent. As the estimate of a commercial market research firm that does not disclose its method, the number itself is to be treated with the caution it calls for, but its direction is notable, because the one speaking here is not an education provider with its own interest in new talent, but the market researcher who is itself declaring the smaller teams the norm for companies, and in the same breath it tells them that the entry level must not be what gets cut to pay for it.⁵³

“Those who rely on the technology in order to cut junior positions end up having to recruit in the most expensive and most contested segment of the labor market.”

How the expected reduction in employment is supposed to come about at all is answered by a survey that is methodologically in an entirely different league from the self-reports of HR managers. For the working paper “Firm Data on AI” (Yotzov, Barrero, Bloom, Bunn, Davis and others, National Bureau of Economic Research, 2026), four central banks and one Australian university surveyed just under 6,000 executives in the United States, the United Kingdom, Germany, and Australia between November 2025 and January 2026 with identically worded questions, whereby the participants were not paid, their identity was verified, and the panels have been validated against the official aggregate figures over ten years. For the past three years, more than nine out of ten of these companies report no effect of artificial intelligence on their employment whatsoever. For the coming three years they expect an average decline of 0.7 percent, and to a subsample of British firms the researchers put the decisive follow-up question, namely where this reduction is supposed to come from. Around two thirds of it is accounted for not by separations from existing employees,

but by fewer new people being hired. That is the mechanism of this chapter, confirmed from the management side and from a source with no vendor interest. The labor market this expectation describes loses no jobs, it stops advertising new ones, and anyone who already holds a job notices nothing of it. Honesty requires noting that these are expectations of executives about their own company and not a measurement, and that the average figure only comes about because the researchers assign numerical values to five answer categories.⁵⁴

The same survey put identical questions to around 3,000 American employees, and the distance between the two answers is perhaps the most revealing number in the entire paper. The employees expect artificial intelligence to raise employment at their own company by half a percent over the coming three years. Their superiors expect a decline for the same company, in the United States of 1.2 percent. On productivity the relationship reverses, because there the employees are the more cautious side at 0.9 percent and management the confident one at 2.3 percent. Anyone looking for the job killer will therefore find it neither in the payrolls nor in the minds of the workforce, but in the expectation of the executive floor, and so far that expectation has no measured counterpart. The same holds on the other side: the confidence of the employees also rests on nothing but an assumption. Two groups in the same company are in all likelihood wrong in opposite directions, and the public debate is only the loud version of this misunderstanding.⁵⁴

“Anyone looking for the job killer will find it neither in the payrolls nor in the workforce, but in the expectation of management, and so far that expectation has no measured result.”

What the payroll data, the firm data, and the analyst warnings measure from the outside, those affected confirm not only in their expectations but also in the advice they pass on to those coming up behind them, and in a form that gives pause. The same sentiment survey by Lenny's Newsletter that supplied the inside view of employees in the chapter on the first force asked its more than 5,000 working respondents for the first time in 2026 whether they would recommend a career in their own profession to someone starting out. More than half would actively advise against it, and the gradient follows exactly the work experience whose beginning this chapter is about. The more senior someone already is, the more likely they still consider this career path worthwhile, founders advise almost neutrally, experienced professionals in the middle of their career advise against it most clearly, and even a third of those who are optimistic about their own career would no longer recommend their own path to anyone. The authors sum up this stance in the formula that the water is fine but you should not come in: the established feel secure themselves because the technology amplifies their experience, yet they have lost confidence that the path by which they came up is still open to those coming after them. It is notable that in its recommendations to executives the survey issues the same warning as Gartner, namely not to let the entry level dry out, because the entry level work that the machine does today was the teaching material on which the experienced of tomorrow learn. These figures too call for the usual qualification, because they come from a self-selected readership of an American industry newsletter, shaped by product managers and predominantly experienced, and they measure attitudes, not actual new hires. Yet that is precisely where their own weight lies, because here it is not those affected complaining about rejection they

have suffered, but the beneficiaries of the shift themselves testifying that the entry level they no longer need has broken away for the others.²⁸

The devaluation of the entry level has a consequence that reaches far beyond the current debate and is almost always overlooked in the excitement about the headline. If the level thins out on which people used to learn the craft, the question arises where the experienced professionals of the day after tomorrow are supposed to come from. The upgrading of the experienced by the technology lasts only as long as enough experienced people come up behind them, and an economy that neglects the training of its next generation is using up a finite resource. The second force is therefore real, but contested in its cause, and the honest position does not consist in resolving the dispute in favor of the simple story, but in enduring it and at the same time showing where the actual danger lies. That this question does not begin at the labor market is shown by the study of 26,811 Chinese schoolchildren evaluated in the chapter on the third force, in which ability was already declining at school while the work handed in kept looking better.⁵⁵

Looking ahead

The career gap will not be closed by one of the three candidate causes winning the scientific dispute, because the entry level thins out in each of the three scenarios. It will be closed by companies deliberately rebuilding the entry level, and the outlines of this rebuild are already recognizable. If the machine takes over the routine groundwork that used to be the teaching material of beginners, then the learning task of the person starting out shifts from execution to oversight, from the first draft to the examining judgment on the machine's draft, and an entry level understood this way can build experience faster than the old one, not more slowly. The German speaking countries in particular have, in their dual vocational training, a long grown institution in which this principle of guided learning on the real case has been practiced for generations and which can be carried over to working with digital team members, while other labor markets first have to invent something comparable.

“When the machine takes over the routine work that used to be the learning task of the entry level, the entry level shifts from execution to oversight, and experience can be built faster that way than on the old path.”

That this rebuild does not remain a mere wish is shown by the first companies already trying it out. The consulting firm BCG (bcg.com), in whose survey of 70 executives more than half name the slower development of junior staff as a risk, reports on the oil company Shell, which has reversed the learning paths of its entry level hires. There the younger ones must first outline a problem themselves, test their assumptions, and present a first analysis of their own before they bring in artificial intelligence at all, which then refines their work instead of taking it off their hands. The idea hits the core of the career gap, because judgment matures only on the work itself, and anyone who spares beginners the laborious groundwork entirely takes from them the practice out of which experience grows in the first place. As the report of a consultancy with its own interest in such

rebuilds, the example is to be read with caution, but it shows vividly what a rebuilt entry level can look like.⁵⁶

At the same time it is foreseeable that the companies cutting their entry level positions today will, in a few years, have to buy in on the most expensive and most contested segment of the market what they failed to train in junior staff. The first companies to have rebuilt the entry level will then have their lead exactly where the gap is opening up today.

The second axis of the career gap runs along gender, in risk as much as in access

The debate about the career gap is conducted almost exclusively as a question of age, and the previous chapter showed why: that is where the measurements are, and that is where the dispute is. The second axis runs along gender, and it deserves a chapter of its own, because it shows up in two different places that are regularly conflated in the public debate. The first place is risk, that is, the question of whether women more often work in occupations whose tasks the technology can take over. The second is access, that is, the question of whether women use the technology where they work as much as men do. There are now solid answers to both questions, and they come out differently. For the German-speaking region this counts twice, because the German labor market, as the German Institute for Economic Research notes, is strongly segregated by gender, and because the occupations in which demographic departures create the greatest additional demand are predominantly women's occupations.

On the risk side the international evidence is clear in its direction and divided on its magnitude. A study by the International Labour Organization published in March 2026, based on harmonized microdata from 84 countries, finds that female-dominated occupations are exposed to generative AI almost twice as often as male-dominated ones, at 29 versus 16 percent, and that the gap is larger still in the highest exposure category, at 16 versus 3 percent, although here too the distinction between exposure and actual displacement described at the outset applies. The cause lies in occupational choice itself and not in the technology, because women are overrepresented in secretarial, reception and clerical occupations with routine, easily codified tasks, while men work more in construction and manufacturing, where activities are harder to automate.⁵⁷ A methodologically independent study by the International Monetary Fund from early 2024, which uses a broader exposure measure not tailored to generative tools alone, confirms the basic direction, namely that in most of the countries examined women work more often in highly exposed occupations than men, but it adds a nuance. Because this exposure among women is split roughly in half between occupations in which the technology tends to complement and those in which it tends to replace, women therefore carry not only a greater risk but at the same time face a greater opportunity, provided that access to the appropriate tools is open to them.⁵⁷ It is precisely this access that has been the sticking point so far, because a study by the OECD from late 2024 does find a broadly similar unadjusted exposure rate between the genders, but documents a gap of 20 percentage points within the same occupations in the actual use of ChatGPT, at the expense of women.⁵⁷

With the American payroll data on which the Stanford measurement of the career gap rests, this finding can for the first time be tested not only against exposure but against actual employment. Brynjolfsson, Chandar and Chen note in the August 2026 version that women at every age work in more heavily exposed occupations than men, in their own sample as well as in the official population survey. Of women aged 22 to 25, according to the official survey, 37.8 percent work in the two most heavily exposed fifths of occupations, compared with 31.4 percent of men the same

age, and the gap persists across all age groups. The employment declines among the youngest are correspondingly steeper for young women than for young men, which the authors explain by the fact that young women are more concentrated in the shrinking, highly exposed occupations. The occupations themselves also differ by gender, because in the most heavily exposed fifth men are concentrated in technical activities such as software development, and in the least exposed fifth women are concentrated in cleaning and health occupations. The authors themselves urge caution, because their sample overstates the exposed occupations, more so for women than for men, and they explicitly describe the gender question as a topic that will have to be watched further in the coming years.⁴

For Germany a dedicated study of the risk has been available since April 2026, and it corrects the obvious expectation. Virginia Sondergeld of the Hiring Lab of the job board Indeed, at the same time a guest researcher at the German Institute for Economic Research (diw.de), has, together with Katharina Wrohlich and Julia Redelings of DIW, linked her own company's transformation index with the gender shares from the microcensus. That the first author works for the data provider needs to be disclosed, but it changes nothing about the link to the official data. The index estimates, for each occupational group, what share of the skills demanded in job ads can be changed wholly or in part by generative AI. There is no linear relationship between the share of women and transformation potential, but an inverted U. At both ends of the distribution, in childcare with more than 90 percent women as in construction with more than 95 percent men, hardly any core skill can be changed by the technology; in childcare fewer than one percent of skills can be changed completely and two thirds only minimally. The occupations with the highest potential, by contrast, software development at just under 84 percent, accounting at a good 75 and marketing at 73 percent, are predominantly mixed. Women work somewhat more often in occupations with higher potential, but the difference lies in the middle of the distribution: in occupations where 60 to 70 percent of skills can be changed, 23 percent of women work against 12 percent of men, while at the ends both genders are equally represented, at the lower end with 20 to 30 percent changeable skills at around 13 percent each, at the upper end with 70 to 80 percent at around 7 percent each. The institute concludes from this that the transformation potential concerns men as much as women, and that the need for further training exists for all. For this report the flip side of the finding is the more revealing one: the female-dominated care and support occupations, in which the demographics chapter locates the greatest additional demand, are the least reachable for the technology, and so precisely where the staffing gap is largest, the relief that the rest of the economy is hoping for fails to arrive.⁵⁸

For the German-speaking region, the finding on access has been available since December 2025 for the first time as an official measurement, and it shows where the gap arises. Eurostat's statistics on ICT use in households, which asked about generative AI for the first time in 2025, find for Germany that 35 percent of men and 30 percent of women between 16 and 74 used generative AI in the three months before the survey. In private life, men use the tools 20 percent more often than women, 30 against 25 percent. At work the figure is 29 percent, 18 against 14 percent. The gap is therefore larger on the job than at home. Among 16 to 24 year olds it reverses: 57 percent of young women had used generative AI, 49 percent of young men. The gender gap, then, is not brought along from school and circles of friends, it arises in the workplace. Austria shows the same pattern

with the same signs, a lead for men of 15 percent in private use and 30 percent at work, and a lead for women of 14 percent among the young.⁶

What lies behind this is shown by a special analysis by the Institute for Employment Research together with the Initiative D21. It rests on a survey from July 2024 and therefore no longer reflects the current size of the gap, but it does reflect its conditions. Where the employer funds the training, the gap among intensive users shrinks to a single percentage point. Where the knowledge travels informally through technology-minded circles of friends and professional communities, by contrast, use rises only among men and the gap widens.⁶ The gender axis of the career gap is therefore less a question of inclination than one of organization, and it is decided by the very factor that the chapter on the long rebuild identifies as the strongest of the entire rollout, paid and explicitly permitted training.

“The gender axis of the career gap is less a question of inclination than one of organization, and it is decided by paid training.”

The international sources contradict one another on the question of how unequally unadjusted exposure is actually distributed, and the German study shows that in this country the axis runs through the middle of the occupational distribution and not along its edges. But they agree precisely on the point that matters in this chapter. Whether through occupational choice or through access to the tools, the burdens and the opportunities of the technology are not distributed at random but follow existing patterns in the division of labor, and anyone who looks only at age sees only one half of the career gap.

Looking ahead

The gender axis will not be decided in the coming years by the technology, but in the same place as the age axis, namely by how companies organize training. According to the IAB analysis of 2024, the German usage gap shrinks to a single percentage point where the employer pays for further training, and it grows where the knowledge travels informally through technology-minded circles of friends. That is an unusually clear instruction for action on a question that is otherwise readily treated as a matter of culture, and it matches what the chapter on the long rebuild identifies as the single strongest factor of the entire rollout. Two longitudinal questions are thus set for the follow-up editions of this report. The first is whether the American observation that employment of young women in exposed occupations declines more steeply than that of young men repeats itself in the German-speaking region once the delay has been made up, or whether dismissal protection dampens this axis as well. The second is whether the gender gap at work of 29 percent that Eurostat reports for Germany in 2025 becomes measurably smaller in the annual figure for 2026, and whether it shrinks first in those companies that pay for the training. The additional demographic demand in the care and support occupations, where the technology can bring hardly any relief, remains the quantity that puts all the others in perspective. It does not shrink because office work becomes more productive, and it helps decide how many women will be able to move at all, in the coming years, into those occupations where the technology would offer them the greater opportunity.

The long rebuild is the most consequential of the three forces, and for now it is failing not because of the machine but because of the organization

Now that the loud first force has turned out to be a capital market phenomenon and the second has proven to be a measurable but contested collapse at the entry level, what remains is the third, the one that has the least in common with the pace of the headlines and is therefore the easiest to overlook. It is the real change that artificial intelligence sets off inside companies, and it does not run through the spectacular replacement of entire workforces but through the laborious reworking of processes, through the question of which task a digital team member takes over reliably and which it does not, and through the even harder question of how an organization builds the knowledge that makes this new division of labor possible in the first place. This force is the most consequential of the three, but it is the slowest, and its slowness is not an omission, it is its actual nature.

Just how large the distance between the promise of the technology and its actual effect inside the company turns out to be shows most clearly in the gap between pilot and production. A widely noted study by the NANDA initiative at MIT from July 2025, which carries the fitting title of the great AI divide, found, on the basis of interviews with senior managers, a survey of executives and an analysis of more than 300 publicly documented AI projects, that 95 percent of the organizations examined got no measurable success out of their initiatives and that only about 5 percent of the pilots already integrated had a noticeable effect on revenue.⁵⁹ The figure is contested in the sharpness of its claim, and it deserves the same caution as any catchy ratio, but its direction is confirmed by independent sources. McKinsey's large stocktaking from November 2025, which surveyed just under 2,000 specialists and executives, comes to the result that only around a third of companies have begun to take artificial intelligence beyond the experiment and into broad use, while two thirds still have that step ahead of them, and that only around 6 percent of respondents attribute to the technology both at least 5 percent of the operating result and a noticeable contribution to value, while not even four in ten see any effect on the result at all.⁶⁰ Two methodologically quite different approaches thus arrive at the same picture, namely that the technology is widely tested but has so far hardly become truly productive anywhere.

Nine months later the continuation of the same survey series is available, and its finding is more revealing for this chapter than any single record number. For the edition of August 25, 2026, a good seventeen hundred participants from 97 countries answered between May 4 and June 8, and adoption has clearly increased this year. 44 percent report that their company is putting the technology into broad use across the entire enterprise, after 38 percent in the previous year, and 56 percent instead of 51 percent use it in three or more functions. The share that attributes any effect at all on the operating result to it still stands at 37 percent and is thus practically unchanged, and the group with at least 5 percent of the result and at the same time a noticeable contribution to value remains at around 6 percent. A year of broader adoption has therefore not narrowed the

distance between pilot and production. Its sharpest form appears in the same survey in the difference between the individual person and their company. 80 percent of respondents report that the technology has improved their own productivity, 50 percent that it helps them make better decisions, yet only 37 percent find any of that again in their company's result. What the individual gains does not arrive in the organization. Where it does arrive, one difference stands out that fits everything this chapter has shown so far. Among the 6 percent with a measurable contribution to operating profit, just under three quarters say they rebuilt their workflows from the ground up because of the technology, in the previous year it was 55 percent, among all other respondents it is 25 percent.³


“Among the 6 percent with a measurable contribution to operating profit, just under three quarters rebuilt their workflows from the ground up because of the technology.”

What sets the small group with a measurable contribution to operating profit apart from all the others is examined in the same survey across eleven operational practices, and the ranking of the differences is more revealing than any single one of them. The two groups are furthest apart on the rebuilding of workflows, closely followed by the ambition to reshape the whole company instead of only cutting costs, and by the question of whether top management visibly takes responsibility for its initiatives itself, which twice as many of the successful companies report. They are closest together in precisely the place where one would most expect the difference, namely on the question of whether the company has set up a dedicated role with the authority to speed up adoption and clear obstacles out of the way. Such roles exist about equally often in both groups. Granting authority is therefore not enough. The goal setting also turns out differently than expected: the successful companies do not pursue the efficiency goal more often than the others but somewhat less often, 78 versus 82 percent. The difference is that they additionally aim at growth and innovation, 74 versus 47 and 65 versus 49 percent.³

“The successful companies do not pursue the efficiency goal more often than the others, the difference is that they additionally aim at growth and innovation.”

What the individual gains barely arrives in the organization.

Share of respondents, in percent



Source: McKinsey, The State of AI in 2026, August 25, 2026, 1,719 respondents in 97 countries.

Note: Self-reported; McKinsey sells AI rebuild programs.

FIGURE 4

What the individual actually gains can by now be quantified more precisely than through self-reports, because for the effect on the single task there are two randomized experiments that count as a reference in the academic literature. Shakked Noy and Whitney Zhang of the Massachusetts Institute of Technology had a total of 453 college-educated professionals work on job-related writing tasks of 20 to 30 minutes in 2023, half of them with access to ChatGPT. The time spent fell by 40 percent, the quality as rated by blinded assessors rose by 18 percent, and inequality among the participants decreased, because those who had been weaker before gained the most. The peer-reviewed version appeared in *Science*.⁶¹ In the same year, Fabrizio Dell'Acqua and Karim Lakhani of Harvard Business School, Ethan Mollick of the Wharton School and their coauthors randomly divided 758 consultants at Boston Consulting Group (bcg.com) into groups with and without access to GPT-4; that was around 7 percent of the consultants without management responsibility at that company. Within the performance frontier of the machine, the consultants with AI completed 12.2 percent more tasks, were 25.1 percent faster and, according to the peer-reviewed version published in *Organization Science* in 2026, delivered quality that was about a third better; three of the authors worked at the company under study, which the study itself discloses.⁹ That is the best available evidence for the gain of the individual, and it shows a time saving of a quarter to two fifths on bounded single tasks of under half an hour, measured with the models of 2023. That is far from a doubling of a person's effectiveness and further still from a doubling of a company's performance. Anyone reading about the 80 percent of McKinsey respondents who report productivity gains of their own has in these two experiments the yardstick for what such a gain means for the individual: a noticeable saving on a partial task that arrives in the operating result only once the organization turns the time gained into other work.

Since the spring of 2026 the distance between pilot and production can moreover, for the first time, not only be asked about but measured on the actual usage data itself. The same Anthropic research note that documents the absent unemployment effect in the chapter on the career gap examines two things for every occupational task, namely whether it could in theory be taken over by language models and whether it is in fact already being taken over in professional use of the vendor's own platform. The distance between the two is enormous. In the most exposed occupational category, the computer and mathematical occupations, 94 percent of tasks would be feasible in theory, 33 percent are actually taken over, and across total employment no use of their activities can be measured at all for three in ten workers. It is also revealing that only this usage measure is related to the independent employment projections of the American labor statisticians, while purely theoretical feasibility on its own shows no relationship whatsoever. It is therefore not what the technology could do that changes work, but that part of it which organizations actually put to use. This confirms the double finding of MIT and McKinsey from an entirely different direction, from usage data instead of surveys, although this measurement too is to be read as the self-observation of a single vendor whose customer base is not a random sample of the economy.⁴⁰

The same vendor supplies a second measurement from usage data with its ongoing Economic Index, which captures the same gap from another side. Since the beginning of 2026 the index has systematically compared how people use the same model in chat and how organizations use it through the programming interface, and in doing so it runs into two separate worlds. Companies use the model three quarters of the time for pure automation, they hand it the task without follow-up questions, and on average the process is over after 5 minutes, while the people in dialogue iterate, adjust and learn. The price of this automation is measurable, because in automated use by companies the estimated success rate is 49 percent, in human dialogue 67, so the task succeeds less often precisely where no one clarifies and corrects any more. If Anthropic weights its own productivity prediction honestly with these success rates, the contribution to American productivity growth calculated from it is halved, from 1.8 to around one percentage point per year. This measurement too is the self-observation of a vendor with all the reservations of the previous paragraph, but on the basis of the usage data it shows what the surveys of MIT and McKinsey report from inside companies. The gap between pilot and production does not arise because the machine can do too little, but because the organization automates away the part of the work on which its success depends, namely the clarifying, adjusting and learning involved in dealing with it.⁶²

Since July 23, 2026 this diagnosis no longer rests on the usage data of a single vendor alone. With the AI & Economy ATLAS, Google and Google DeepMind have for the first time measured the use of their own services, almost 15 million anonymized interactions across the Gemini app, the AI mode of search and the programming interface from two weeks of April 2026, accompanied externally by economists such as David Autor and Diane Coyle. The result repeats the pattern with different numbers: Gemini use is observable in 68 percent of all occupations, which together account for a good 88 percent of American employment, yet at the median it touches only 21 percent of the tasks of an occupation, and in only 3 percent of occupations does it reach more than three quarters of the tasks. The direction of use matches as well, because in demanding knowledge work fewer than one in ten conversations aims at handing a task over to the machine completely, the bulk is preparatory work, revision, idea generation and looking things up. The

technology is therefore widely tested but used shallowly, and it predominantly accompanies the work instead of replacing it: what the usage data from Claude showed, the data from Gemini shows just the same.⁶³

With OpenAI (openai.com), at the end of July 2026 the third large model vendor also presented an ongoing measurement of the use of its own services, and it adds to the picture the question of whose work people are actually handing to the machine. For the first edition of the series *Work at the Frontier*, OpenAI mapped more than 800,000 work-related messages from American ChatGPT users to the activity catalogs of the official occupational database O*NET. A good 60 percent fall to general work such as emails and summaries, yet of the requests that can be assigned to a discipline, 43.5 percent concern tasks that traditionally belong to another occupation: whoever works in customer service, in design or in human resources predominantly pulls in work from outside their own field on the specialist requests, runs the numbers on financial data, fixes software problems or drafts advertising material. The pattern is most pronounced in the smallest units, because among typical users the out-of-field share in workspaces with two to five licenses is around an eighth higher than in those with more than 100; there, the person who reaches for the machine is the one who has no specialist beside them to delegate to. OpenAI itself draws a remarkably cautious conclusion from this: the technology changes not only how work is done but who does it, while the specialists remain indispensable for judgment and review. This measurement too is the self-observation of a vendor at the level of individual messages, without information on the quality and use of the results and, unlike at Anthropic and Google, without a disclosed analysis period. But it adds the missing dimension to the series of usage data: while use per occupation remains shallow, the tasks are already migrating across the organization. Occupational boundaries are shifting faster than the work is passing to the machine, and this shift too is not an achievement of the technology alone, but a question of how an organization arranges review, responsibility and collaboration around the new possibilities.¹⁰

The same gap between technical feasibility and actual use can also be documented independently of the vendors themselves, and officially so for Germany. The substitution potential of the Institute for Employment Research has measured at the task level since 2013 what share of the core activities of an occupation could already be performed technically by computers or computer-controlled machines, explicitly as a potential and not as a replacement that has taken place. In 2022, 38 percent of employees subject to social security contributions worked in an occupation with a high substitution potential, compared with 34 percent in 2019, yet how far this potential is actually used remains another question, precisely that distinction on which the Anthropic figures and the gap between pilot and production in this chapter also rest. Just how movable this yardstick becomes visible in a case from the measurement itself. The training occupation of tax clerk was considered 100 percent automatable by the institute from 2016 onward, because for each of its typical activities a technology could be found that could have taken it over completely. With the new training regulation that came into force in 2023, the same value fell to 50 percent, because the occupation had been described anew, with more support, advice and strategic tasks and less pure bookkeeping. The technology had not fallen behind, the occupation had shifted, and the measurement followed. Occupations rarely disappear entirely, they are rewritten, and whoever reads the automatability of an occupation as a number is always also measuring how old its description is.⁶⁴

What matters is why that is so, because at this point the diagnosis shifts away from the machine. The MIT study explicitly does not attribute the failure to the quality of the models but to a learning gap of the organizations, that is, to the difficulty of fitting the technology cleanly into existing processes and of building the knowledge required for that inside the company.⁵⁹ That is the real bottleneck of the third force. The models are ready faster than the organizations that are supposed to absorb them, and the bottleneck lies not in the data center but in the ability of a company to learn how a capable tool becomes a productive digital colleague.

“The bottleneck lies not in the data center but in the ability of a company to learn how a capable tool becomes a productive digital colleague.”

For the distinction between a technology that supports existing work and a technology that organizes work from the ground up, a study by Hyunjin Kim and Rembrand Koning on American AI start-ups supplies a sharper pair of terms, even though it starts from quite different firms than the established corporations of this chapter. It separates a process channel, in which AI tools such as a language model or a programming assistant support the work of existing employees, from a product channel, in which AI capabilities are built directly into the product that is sold, so that the knowledge work migrates into the product itself instead of into an internal team. In their sample, only the product channel predicts the size and the depth of the hierarchy of a firm, while the mere use of AI tools inside the company, widespread as it is, significantly influences neither the one nor the other.⁵¹ Applied to the third force discussed here, this means that the real starting point does not lie in equipping employees with one more tool, but in rebuilding processes and products from the ground up around the capabilities of the technology, exactly that act of learning at which the MIT study sees the organizations failing so far.⁵⁹

That this finding is not an American peculiarity but describes the German economy across its full breadth is shown by a study by IW Consult commissioned by eco, the association of the internet industry, from July 2026, for which 500 companies weighted by size were surveyed. It determines that around 40 percent of the German economy now use artificial intelligence, an increase of 118 percent since 2024 that reaches down into small businesses with fewer than 50 employees, and it captures the German approach in the fitting formula of refinement instead of reinvention. Instead of putting billions into developing their own foundation models, companies take globally available models and adapt them to their specialized processes with their own industry data, that is, exactly that work of adaptation which keeps the knowledge inside the company instead of giving it away to an outside product. More revealing than the boom figure itself, however, is how narrow the top of this movement remains, because only about one company in five already uses the technology actively for innovation, and among the users 46 percent reach exclusively for finished standard solutions, while merely 10 percent really develop the models further in a targeted way and another 44 percent combine both routes. The breadth of use and the depth of its productive anchoring therefore diverge clearly in the domestic figures as well, which confirms the gap between pilot and production once more, this time in the statistics of the German economy itself. That the survey was commissioned by an industry association which explicitly derives from it the call for less regulation

belongs to the assessment, but it changes little about the basic direction, because it is consistent with the observations of MIT, of McKinsey and of the Bundesbank.⁶⁵

How independent this organizational decision is of the mere assessment of the technology is shown by a more recent study by Innessa Colaiacovo and Rembrand Koning, the same coauthor as in the study just cited, among just under 200 founders of North American technology start-ups, one of the groups most familiar with artificial intelligence of all, since nine in ten use it personally. By means of a randomized experiment the authors exposed the founders to optimistic or pessimistic information about the capability of the technology in customer service and showed that while the assessment of the technology can be shifted measurably, the willingness to actually hand it work is explained by that, even afterward, only to a good third. Founders who formally integrate artificial intelligence into their workflows differ from those who merely leave it to their staff in passing hardly at all in their assessment of the technology itself, but clearly in their actual delegation behavior and in the staffing effect they attribute to the technology: without it, the one group estimates it would need 91 percent more staff, the other 24 percent. Even in this particularly technology-minded sample, just under a third of the founders moreover state that they have realized no staff savings at all, which confirms the gap between pilot and production from the side of the users themselves and at the same time sharpens it, because by this account the bottleneck lies neither in access to the technology nor in belief in it, but in the independent, laborious decision to actually redivide the work, as the authors themselves sum it up when they write that the effect of artificial intelligence on work and organizations depends not only on what the technology can do, but also on what executives are prepared to do with it.⁷

That this diagnosis is no peculiarity of management consulting but reaches all the way into the central banks is shown by the keynote address “Already Here, Not Yet Everywhere,” which German Bundesbank President Joachim Nagel delivered on April 21, 2026, at the International Economic Symposium in Rome, and which captures the whole matter with a phrase from the writer William Gibson, according to which the future is already here, just very unevenly distributed. Applied to artificial intelligence, this means that its capabilities have long been visible in many applications while its macroeconomic effect still barely appears in the aggregate statistics, and for this delay Nagel reaches for the same historical parallel that goes to the core of the third force. When electricity spread through the economies, the productivity gains at first remained small, because the plants did install the electric motors but did not reorganize their production right away, and it was precisely this reorganization, not the availability of the machine, that decided success in the end.⁶⁶ The most famous example of this reorganization is Henry Ford's assembly line of 1913. The Highland Park plant was designed from the ground up around electric power available everywhere, because only that made it possible to arrange the machines in the order of the work steps and to drive the moving belt itself, whereas a classic steam factory had to group its machines around the central transmission shaft and its belts. Economic historians such as Warren Devine and David Nye record that mass production of the Model T would not have been practicable without electricity.⁶⁷ Volume production of the automobile was therefore not a consequence of the new machine alone, but of the newly ordered work around it. It is the same learning gap that MIT observes at the project level, now seen from the altitude of monetary policy.⁵⁹

“It is not the availability of AI that decides success, but the reorganization of work around it.”

For the German-speaking countries the same Nagel speech supplies the concrete confirmation, because it draws on the Bundesbank's own company survey. By that account most German firms do expect productivity gains from generative artificial intelligence, yet the intensity of actual use continues to rise only at a moderate pace, which suggests that many companies are still experimenting rather than moving the technology into day-to-day operations. The accompanying data in the Bundesbank monthly report of March 2026 put figures to this picture, because while the share of German firms that use generative AI tools is high by European comparison, the share of working time in which they are really put to use grows only slowly from year to year. The gap between pilot and production thus no longer appears only in American and global surveys, but in the figures of the German economy itself, and the third force loses the appearance of a distant phenomenon that concerns only the front pages of Silicon Valley.⁶⁶

“While the share of German firms that use generative AI tools is high by European comparison, the share of working time in which they are really put to use grows only slowly from year to year.”

Just how concretely this gap shows up in a single, particularly regulated industry is documented by a study from Wavestone (wavestone.com) and the Institute of Information Systems at the University of St. Gallen, which between October 2025 and May 2026 surveyed a total of 30 business and IT leaders from banks, insurers and technology vendors in the German-speaking countries. The finding is the sharpest in this chapter so far, and it even has a face. The most advanced interviewee heads the central AI lab of a large insurer, already runs 75 productive use cases there in classic machine learning and generative AI, among them an application that searches some half a million court rulings, and has still not brought a single agentic use case into production. He himself rates his institution at only a three on a maturity scale of ten, and in that shows the distance between an institution that operates AI and one that lets itself be changed by it. The study sums up the result itself in one sentence that confirms this report's central thesis from an entirely different industry: it is not the capability of the models but the maturity of the operating model that decides how far a company can actually scale agentic AI. Regulation sharpens this further, as the AI lead of a German bank discloses: “We process personal and highly sensitive data here. Technically that is not the problem, but the regulatory effort, the AI Act, the GDPR, DORA, those are the biggest constraints.” As an organizational answer, a central-decentral model is winning out at most of the institutions surveyed, a lean central AI unit for governance and shared capabilities, coupled with business-specific use cases in the individual divisions. Part of the context is that Wavestone, as a consulting firm, takes its share of exactly the transformation it examines here, while the academic involvement of the University of St. Gallen secures the study's independence from individual vendors.⁶⁸

“The insurer runs 75 productive use cases, among them an application that searches half a million court rulings, and still has not a single agentic use case in production.”

The same gap between trial and operation shows up across the breadth of small and medium sized enterprises, for which there are only a few measurement series in the German-speaking region. Since 2024 the Swiss SME labor market study by AXA and Sotomo has asked not only whether companies with five to 250 employees use artificial intelligence, but how. In March 2026, 35 percent integrated the technology deliberately into their processes, 13 percent in all areas of the company, 39 percent were trying it out in individual projects or on the employees' own initiative, and 26 percent did not use it. Trial and use together jumped from 55 percent in 2024 to 71 percent in 2025 and stand at 74 in 2026; the jump is over, the plateau is reached. More revealing is the series of fields of use. In five of eight fields the 2026 value lies below that of 2025: translations fall from 52 to 47 percent, correspondence from 47 to 42, personalized advertising from 24 to 16, while the optimization of work steps holds at 35 percent and data analysis keeps growing from 32 to 34 percent. The text tasks with which the first use began reached their peak in 2025; what remains are the applications in the workflows. The study itself compares only 2024 with 2026 and calls that stagnation, and so the peak of 2025 disappears from its own presentation. The size axis runs continuously all the way down to the small firm: 23 percent of companies with five to nine employees see a need for action on their own business model, 38 percent of medium sized ones and 42 percent of firms with 50 or more employees, and 33 percent of the small ones rate the technology as an opportunity against 49 percent of the medium sized ones. Behind each size class stand only around 100 respondents, and micro firms with fewer than five employees are missing entirely.¹³⁵

“With personal data the problem is not the technology but the regulatory effort from the AI Act, the GDPR and DORA, says the AI lead of a German bank.”

That this gap can be closed, and what it actually costs an organization to close it, has been documented since July 2026 by, of all things, a government agency. The US State Department, with its Generative AI Playbook, has laid open the full path of its internal assistant StateChat, from the first data strategy to operation with more than 62,000 users at 98 percent of its diplomatic posts worldwide, out of a workforce of some 80,000 people. What is remarkable about this document is less the success than the accounting of the effort behind it. From the build-out of its own analytics center in 2019 through the data strategy of 2021 to the preparation of agentic systems in the spring of 2026, seven years passed, the security clearance for sensitive data alone cost another year of procedural work, and the department describes precisely that not as a brake but as a precondition of the benefit, because without access to the real working data the tool would have remained worthless, a remarkable counterposition to the complaint about regulatory effort that the previous paragraph quotes from a German bank. The playbook devotes only a fraction of its pages to the technology; the rest is given over to adoption in everyday work, 94 open office hours with some 17,000 participants, more than 430 training sessions with more than 48,000 people trained on the

train-the-trainer principle, and, as the single most effective measure, explicit permission from above, because the most common inhibition among staff was uncertainty about whether they were allowed to use the tool for their work at all. After the leadership of one regional bureau personally invited their people, the number of weekly active users there rose by 367 percent.⁶⁹

The successes the department measures after a year of operation are modest, and credible for exactly that reason. Nine out of ten regular users report time savings, on average 1.6 hours per week and task, most often when drafting emails and memos, and among those who gain no time the internal survey names as one main reason that checking the machine's results takes longer than doing the work themselves would have taken, that is, the supervisory work this chapter will come back to, documented here for the first time from inside an organization with 80,000 employees. All of this is to be read as the self-presentation of an agency that wants to write itself to the front of the American administration, based on a voluntary survey without independent productivity measurement, and what becomes of the freed-up hours structurally is left open by the document. But as a publicly viewable record of what the long rebuild really costs a single organization, it is so far without counterpart, and its finding agrees with every measurement in this chapter, because at no point was the bottleneck the model, but the organization that first had to be built around the model.⁶⁹

A global survey by the software vendor Notion (notion.com), which in the spring of 2026 covered more than 6,000 decision makers and users in 16 markets, confirms the same diagnosis once more, this time not on the profit and loss statement but on the maturity of the organizations. 88 percent of organizations still use the technology where it serves as an individual tool for drafts, summaries and first ideas, and only 12 percent have already built it into recurring workflows or let it run autonomously. More telling than this distribution is what changes as maturity grows. The share of decision makers who concede that their organization is investing in the technology faster than its own workforce can keep up with it grows with increasing maturity from just under half to more than two thirds. The learning gap therefore does not close the further an organization advances; it opens wider, and thus it is confirmed from a third, independent side that the real bottleneck does not lie in the availability of the technology but in the pace at which an organization builds the knowledge needed to use it. Part of the context is that Notion itself sells AI-supported work software and has its own interest in presenting deeper system integration as the decisive leap in maturity, but the direction of the finding agrees independently with the MIT and the McKinsey figure as well as with the Bundesbank observation.⁷⁰

“The decision makers confirm the gap not on the profit and loss statement, but in the maturity of their own organization.”

“88 percent of organizations use the technology only as a tool for drafts, summaries and first ideas, only 12 percent have built it into recurring workflows.”

The more mature the organization, the more often learning lags behind.

Decision makers whose company invests faster than the workforce learns, in percent



Source: Notion, The Great Renovation, spring 2026, 2,039 decision makers in 16 markets.

Note: Notion sells AI software itself.

FIGURE 5

How this gap is by now arriving in the executive suites themselves is shown by the third annual AI at Work survey from the employment services provider G-P (globalization-partners.com), for which the field institute Wakefield Research (wakefieldresearch.com) surveyed a total of 2,850 executives from vice president level upward in six markets in March 2026, among them Germany. Without exception, all respondents now use the technology, yet 73 percent concede that the success of their AI spending has fallen short of expectations, the share of those who describe their own approach as an offensive innovation strategy fell within one year from 60 to 42 percent, and just under 70 percent announce that they will cut AI budgets should the company targets for 2026 be missed, in Germany at 58 percent the most reserved figure, but still a majority. What is new about this year's edition is therefore less the diagnosis, which agrees with all the findings gathered so far, than the deadline it sets for it, because the patience that the long rebuild by its nature demands is being limited, where the budgets are approved, to a single year, and with that the slowest of the three forces comes under a beat that contradicts its nature. That these are self-reports from a vendor survey belongs to the context, because G-P, as a provider that formally employs staff on behalf of other companies, sells cross-border employment together with its own AI agent and advertises both in the same report, yet precisely because this business lives on belief in the technology, the documented disillusionment weighs more heavily rather than less. For the question of this chapter that means the gap between pilot and production has arrived in the profit and loss statement and is no longer read there as a maturing phenomenon that is given time, but as a bill that has to be settled within a single fiscal year.⁷¹

“Just under 70 percent of executives announce that they will cut AI budgets should the company targets for 2026 be missed.”

“The gap between pilot and production has arrived in the profit and loss statement, and there it is no longer read as a maturing phase, but as a bill that has to be settled within a single fiscal year.”

How this impatience squares with the actual movement of companies is shown by the most recent large-scale measurement to date, the second edition of the quarterly Global AI Pulse from the consulting firm KPMG (kpmg.com), which between the end of April and the end of May 2026 surveyed a total of 2,145 executives in 20 countries. The share of companies that anchor the technology broadly in their own operations has nearly doubled within a single quarter, from 13 to 22 percent, the largest movement the survey has measured anywhere on its maturity curve so far, and at the same time the share of those that can show an established, demonstrable success fell slightly from 8 to 7 percent. The gap between pilot and production therefore widens at exactly the moment when three quarters of respondents are convinced that the technology already delivers measurable value to their company, and 79 percent state that they would hold on to that priority even in a recession. The companies' answer to this tension is not retreat but management, because the planned budgets stay stable while almost half of respondents have already re-timed the rollout of their AI agents as soon as the costs began to exceed the expected benefit. More telling than the gap itself is what distinguishes the companies with demonstrated success from the rest, because it is not the level of spending but the order behind it. Where responsibility for the results of the technology is clearly named, an established success is reported more than three times as often, 14 versus 4 percent, and where an organization has full visibility of its AI running costs, five times as often, 15 versus 3 percent, while three quarters do say their CEO drives the topic as a strategic priority, but only a quarter actually locate ultimate responsibility at the top. That these figures come from a consulting firm that earns its money closing exactly this gap belongs to the context, but the finding adds something to this chapter's diagnosis that the earlier snapshots could not provide, namely the proof over time that the gap is not a question of the technology or of the money, but of the organization, which has yet to move responsibility and cost visibility into its structures.⁷²

“The gap between pilot and production widens at exactly the moment when three quarters of respondents are convinced that the technology already delivers measurable value.”

What this order looks like in operation was publicly demonstrated in the summer of 2026 by a corporation that three months earlier was still the counterexample. The ride-hailing company Uber (uber.com) had used up its entire AI budget for 2026 after just four months, and President and Chief Operating Officer Andrew Macdonald conceded in May that the connection between this spending and the features actually shipped could not be established so far.⁷³ In the quarterly report of August 5, 2026, Chief Financial Officer Balaji Krishnamurthy then reported that the cost per token had fallen over the past months while usage had continued to increase, so that total spending on artificial intelligence remained largely stable.⁷⁴ Chief Technology Officer Praveen Neppalli Naga disclosed on the same day how that had been achieved, and none of it concerns the machine: a higher hit rate in the cache for recurring requests, changed defaults for model choice and context

length, the routing of simpler tasks to cheaper and open models and, most tellingly, a display that shows every developer their own AI costs per hour in real time. Efficiency, as he explained it, had been treated as an engineering task and not as a budget question.⁷⁵ What tipped here within a single year at the same company was therefore not the technology and not the size of the budget, but cost visibility, and at this point the digital workforce visibly gets, for the first time, what every human one has long had, an hourly rate.

What this case does not deliver belongs to the context, and it is exactly the quantity Macdonald was concerned with in May. What fell was the price per unit, and it fell in the market over the same period as well, because the vendors cut their list prices and open models moved up, so that from the outside it cannot be separated which part of the decline goes back to the company's own engineering work and which to the market; the company names no amounts. What is answered is the question of the costs, not the question of the success. It is telling which metric the chief financial officer cites for the benefit, namely a doubling of code output per developer, precisely the quantity this report rejects in the chapter on the technological substructure as evidence of output and not of value.⁷⁶ Of the three tasks that belong to the cost accounting of a digital workforce, measuring the costs, attributing the benefit and proving the contribution to operating profit, the first is thus solved and the other two are open. That is no small step, because without the first the other two cannot be had, but it is not the proof of benefit it is currently taken for.

How far the distance between technical maturity and operational effect reaches can be tested most precisely where the replacement was announced earliest and loudest. In 2016 Geoffrey Hinton recommended stopping the training of radiologists, because within five years it would be obvious that machine learning reads images better than humans do. Ten years later the official French health statistics keep the books on what came of it. The number of radiologists and nuclear medicine physicians in France rose between 2014 and 2024 from 9,100 to 10,200, an increase of 13 percent, that of radiology technologists from 29,000 to 32,000, and for their recruitment the same statistics explicitly record shortages for 2024. The number of outpatient imaging centers grew over the same period from 860 to 1,070, spending on their services rose in 2024 by 7 percent to 5.5 billion euros, carried by volume growth of 8 percent with falling prices. As the cause of this demand the agency names not the technology but the aging of the population, and artificial intelligence appears in the same account exclusively as an expectation for the coming years, as one of those medical fields in which it could unfold and bring productivity gains. In the discipline, then, whose data suit the machine best, the replacement has not occurred a decade after it was announced, and the official statistics still carry it in the conditional.⁷⁷

Where the technology actually arrives in operations, it so far relieves rather than replaces. A before-and-after study by the Charité in Berlin in German long-term care shows that an AI voice assistant cuts documentation time per early shift by 15 minutes, or some 28 percent. The technology therefore takes hold exactly where the burden is greatest, without making a single care worker redundant; the study, however, has no control group, and its lead author works for the maker of the product under study.⁷⁸ At the same time, use is distributed highly unevenly. In Anthropic's usage data, some one third of workers account for no measurable use at all,⁴⁰ and even where Google's analysis of Gemini usage observes the technology reaching into the workshop, for instance with mechanics who interpret test results and upload images of faults, it captures only a narrow slice of the occupation.⁶³ The third force therefore works narrowly so far,

more augmenting than replacing and very unevenly, and that is almost the opposite of the across-the-board displacement the headline claims.

Relief on the individual task has a flip side in the working day, and there the workforces pay for their organizations' learning gap

How fragmented the working day has become was measured by Microsoft in June 2025 on the basis of the anonymized usage data of its office software, which gave the finding the telling name of the infinite workday. By that account an employee is interrupted during core working hours on average every 2 minutes by a meeting, an email or a chat message, for the most heavily burdened fifth of users this adds up to some 275 interruptions a day, meetings after 8 p.m. increased by 16 percent within one year, and in the accompanying survey of 31,000 knowledge workers in 31 countries almost half of the employees and more than half of the managers describe their own work as chaotic and fragmented, while a third say the pace of work of the last five years is simply no longer manageable. With its AI assistants Microsoft sells at the same time the remedy meant to lead out of this situation, so the figures are to be read as the diagnosis of a vendor with its own interest, yet the most remarkable sentence of the report turns against its own sales logic, because it warns that without a newly ordered rhythm of work, artificial intelligence threatens merely to accelerate a broken system.⁷⁹

“During core working hours an employee is interrupted on average every 2 minutes, and without a newly ordered working day artificial intelligence only speeds up this broken system.”

Why this intensification appears in no productivity statistic is explained by a laboratory experiment that the computer scientist Gloria Mark, the leading researcher on the interruption of knowledge work, published together with colleagues from Humboldt University as early as 2008. Test subjects who were interrupted constantly did not complete their tasks more slowly or worse, but faster, because people compensate for expected disturbances with a higher working pace, and they paid for it with significantly more stress, frustration, time pressure and effort, measurable after just 20 minutes of interrupted work. In business terms the calculation therefore appears to work out, while its costs migrate entirely into the people, and that is exactly why the intensification stays invisible until it becomes visible as exhaustion.⁷⁹

“The intensification stays invisible until it becomes visible as exhaustion.”

The long rebuild adds a new layer to this old mechanism, because the digital team member does not only take work off people's hands, it also demands supervision. Anyone who directs several AI agents in parallel, checks their interim results, corrects their misfires and coordinates their handovers is no longer working at their own rhythm but at the rhythm of the machines, and Gartner describes precisely this role in its picture of small teams, where every team member carries, alongside their own specialty, the supervision of AI agents.⁵³

“The digital team member does not only take work off people's hands, it also demands supervision.”

How widely this supervisory work has already arrived is shown by the G-P survey already cited on the budget caveat, in which 69 percent of executives state that they now spend more time monitoring unchecked AI work and reworking it like editors, which is why the report simply declares 2026 the year of the human editor.⁷¹ Orchestration thus becomes the new form of multitasking, and it meets workforces that are in part already working at their limit. The sentiment reading from the survey of tech employees cited in the chapter on the first force fits in here, because the share of those at least moderately burned out jumped within a single year from 44.7 to 55.7 percent, and what the respondents fear is not replacement but the pace, the worry of landing in a rhythm that cannot be sustained, while every gain in productivity immediately becomes the new benchmark of expectation.²⁸ For the diagnosis of this chapter that means that the learning gap is not only a gap in success, but a gap in strain. An organization that invests faster than it learns does not pass the difference on to the machine but to its people, and whether the time freed up is translated into relief or into intensified expectation is the same leadership decision that this chapter already knows from the delegation question.

“Orchestration becomes the new form of multitasking, and it meets workforces that are in part already working at their limit.”

Alongside the gap in strain a second one opens, and it hits the capabilities of the workforce directly. So far this chapter has described the learning gap as the failure to build new capabilities in time. Now the reverse case is added, that a company gradually loses capabilities it has long possessed. The consulting firm BCG calls this distributed skill loss and means by it not that a single person is slipping, but that an entire workforce loses its judgment at the same time, as soon as everyone works with the same tools and on the same data basis. What is revealing is which capabilities the executives surveyed see as most at risk, because they are not the routine steps that are easy to hand over anyway, but the most demanding ones, the sense of which problem is to be solved in the first place, and the judgment of how to decide, that is, precisely those capabilities that a machine imitates effortlessly and truly replaces only with difficulty. The path there is inconspicuous and works for exactly that reason, because since the digital team member delivers usable results quickly enough, one gets used to accepting them as good enough and skips the step that used to be the actual work, the examination of the reasoning behind them. Anyone who stops questioning unlearns it, the way a muscle wastes away when it is no longer needed.⁵⁶

“Now the reverse case is added, that a company gradually loses capabilities it has long possessed.”

“Anyone who stops questioning unlearns it, like a muscle that is no longer needed.”

This finding has to be read with care, because it rests on the assessment of 70 executives and comes from a company that makes its consulting business out of this very problem, and the research on so-called cognitive debt, on thinking weakened by outsourcing, is also still in its early stages and is contested.⁵⁶ Since the beginning of the year, however, the first larger experimental measurement of it has been available. Two behavioral researchers at the Wharton School, Steven Shaw and Gideon Nave, had 1,372 test subjects solve reasoning tasks in three experiments and left it to them whether to consult an AI, nearly 10,000 decisions in total. More than every second task was delegated to the machine, and the direction of the effect depended entirely on whether the machine was right. With correct AI answers the hit rate rose by 25 percentage points, with wrong ones it fell by 15 points below the level of those who worked without AI at all. The most troubling part was the self-confidence, because it grew in both cases, with correct as well as with wrong answers from the machine. The researchers call this reflex cognitive surrender, the adoption of a machine judgment without one's own examination, and they observe it the more often the greater the trust in the technology is, while pronounced analytical thinking protects against it.⁸ The BCG finding itself nevertheless carries little weight as a measurement, and all the more as the self-observation of an occupational group,⁵⁶ because the slackening judgment does not stay with the individual. Where everyone works with the same systems, the results converge, similar answers, similar phrasing, similar conclusions, and the more similar they become, the less reason remains to ask a colleague's opinion at all, while the pace that demands everything at once takes the last of the time away from discussion. At the same time responsibility shifts, because if a result comes from artificial intelligence, it becomes easier not to regard it as one's own any more, and even experienced people replace their judgment with that of the machine out of habit and under time pressure, without anyone noticing the slippage in them. For this chapter it follows that the capability of a company is nothing fixed that one has once and then keeps, but something that persists only in use, and whether the technology sharpens it or hollows it out does not depend on the technology itself but on the same leadership that already decides about delegation and intensification.

“The most troubling part was that the self-confidence grew in both cases, with correct as well as with wrong AI answers. Cognitive surrender is what the researchers call the adoption of an AI judgment without one's own examination.”

That this capitulation happens not only in the laboratory but in everyday working life was shown by the field experiment at the Boston Consulting Group that this chapter cited at the outset for the gain inside the performance frontier, because the same experiment contained a second task deliberately placed outside that frontier. The consultants were to combine a data analysis with qualitative evidence from interviews, and the machine answered plausibly but wrongly. The group with AI access delivered a correct solution 19 percentage points less often than the group without. The same machine that had made the consultants a quarter faster on the first task made them worse than without the tool on the second, and from the outside there was no telling on which side of the frontier the two tasks lay. The authors speak of a jagged performance frontier of the technology.⁹ For the long rebuild this is the experimental evidence for the skipped step of examination. The loss does not arise because the machine is unusable, but because the person stops checking as soon as it has been of use to them once. An organization that books the gain inside the frontier buys the

loss beyond the frontier along with it, as long as it does not teach its people where the frontier runs, and it shifts with every new model.

That this mechanism concerns not only the individual decision but the building of capability itself was shown by a preregistered field experiment in the year before, and its design brings the decisive question to the point. Hamsa Bastani, Osbert Bastani and their team at the University of Pennsylvania had nearly 1,000 students in grades nine through eleven at a large high school in Turkey practice mathematics in four sessions. Every session followed the same sequence, a review by the teacher, a practice phase and a concluding test without any aids, and the only thing varied was the practice phase. One group received a tool that behaved like an ordinary chatbot, a second the same model with an instruction written by teachers that withheld solutions and gave graduated hints instead, a third worked with the usual materials. In the practice both groups with machine assistance were far ahead, the first by 48 and the second by 127 percent. In the test without aids the first group fell 17 percent below the control group, while the second was statistically indistinguishable from it. The same model thus led to opposite results depending on how it was used, and the conversation transcripts show why, because those who could get the solution asked for it and copied it down without making an attempt of their own. For a company the second finding may weigh more heavily than the first, because neither of the two groups recognized its own situation. The one that performed worse did not consider itself worse, and the other, which did not perform better at all, considered itself distinctly better.⁸⁰

What laboratory and self-report can only suggest was tested for the first time in June 2026 by a measurement of ordinary use over years, and it comes not from paid work but from school. David Strömberg of the University of Stockholm together with Victor Lei and Yanhui Wu of the University of Hong Kong followed 26,811 Chinese students in grades seven through twelve over 30 months and made use of the fact that generative artificial intelligence arrived at the participating schools at staggered times, which makes it possible to compare groups with and without access. Homework performance rose by 18 percent, while the time spent on it fell by 30 percent. The results of the monthly tests without aids fell by 20 percent within half a year, and in the large entrance examinations on which the further educational path depends in China, the losses came to 18 and 24 percent, with the full effect only showing after about two years. What is decisive for this chapter, however, is not the size of the loss but its distribution, because it sits almost entirely with that group of around 80 percent of users whose strikingly short working time combined with high homework marks points to an outsourcing of the work rather than to working with the tool. Those who kept up their working time and used the machine alongside it lost almost nothing. It is not the tool that produces the loss, but the giving up of the path to the result.⁵⁵

Transferring this to the company is interpretation and not a measured result, because what was surveyed was a school system in a single Chinese county, and the work is available as a discussion paper and has not been peer reviewed. The mechanism, however, is the same one that this chapter describes for the workforce, and the school supplies what the company has so far lacked, namely a test without aids. A company measures its AI success by what in this study is the homework mark, that is, by the finished result and the time needed for it, and both quantities pointed upward here over months while the underlying capability was already giving way. Anyone who takes this delay seriously reads their own adoption figures differently, because a company

does not get its examination day monthly, but only when a case arises for which the tool has no template.⁵⁵

The BCG finding and the Wharton School experiment describe what happens to human judgment after the machine has answered. How much is already decided beforehand, namely with the question itself, has been measured cleanly for the first time by a not yet peer reviewed paper by Taha Choukmane, Tim de Silva, Weidong Lin and Matthew Akuzawa of the Massachusetts Institute of Technology and Stanford University, albeit in an entirely different field, namely private financial advice. A thousand representatively selected Americans wrote their own questions there to an AI adviser, and the researchers calculated the advice given through an entire working life. The machine's advice was by and large good, it would have brought most respondents closer to what the textbooks recommend. But it turned out differently depending on the person asking, and in order to clarify why, the researchers randomly inserted into those questions that do not state the gender of their own accord the sentence that one is a man or a woman. This made it possible to separate what goes back to the question from what goes back to the model: two thirds of the difference in the recommended equity share arose from people asking differently, one third from the machine answering the same question differently.⁸¹

The finding comes from the financial planning of private households and not from paid work, and transferring it to the company is therefore interpretation and not a measured result. It is nevertheless close at hand, because what the authors record at the end of their paper is the diagnosis of this chapter in different language: the demand side will limit the quality of what people receive even as the models get better. A company that gives everyone the same license and then looks at very different results does not have a model problem, it is watching its own learning gap at work.

The official business survey by the Census Bureau, which in the chapter on the AI argument makes the employment effect come out so small, also measures the depth of the anchoring, and it thereby quantifies the distance between pilot and production more precisely than was previously possible. 57 percent of the companies that use artificial intelligence do so in at most three of their business functions, 65 percent limit their employees' use to at most three types of task, and a typology of users shows that with 37 percent the largest group consists of minimalists, while the technology is used comprehensively across all functions in a full 4 percent of cases, that is, in barely four out of a hundred companies. Even more revealing than this narrowness is how unevenly timed the adoption is. In 36 percent of the companies in which employees work with artificial intelligence there is no formal adoption at the company level at all, so the use has grown from below and past the leadership. Conversely, 19 percent of companies with a formal adoption show no discernible use in their employees' tasks, there the decision has been taken and has got stuck in the implementation. Seeing both at once is the real gain of this survey, because on that basis the spread within a company is not a single uniform process, but consists of two movements independent of one another that can fail to meet.⁸²

The findings of this chapter keep showing the same picture. The distance between what technology is bought and tried out and what of it actually works in daily operation shows up in the official statistics as well as in the surveys and in the vendors' usage data. Precisely this distance between pilot and production is the quantity that the report will measure year after year in a metric

of its own from the 2027 edition on, because it says more about the true state of things than any figure about models or investment sums. How an organization closes this distance, that is, how it builds the knowledge that lifts artificial intelligence out of the trial and into ongoing operation, is the question this report comes down to in the end. At this point the finding suffices that the most consequential of the three forces is so far not failing because of the technology, but because of the organizations that still have to learn how to work with it.

Looking ahead

The long rebuild has a historical script, and its decisive scene is still to come. With electrification the jump in productivity did not come when the factories replaced their steam engines, but only when a new generation of plants was built around the electric motor, with different floor plans, different processes and a different division of labor. The economic historian Paul David of Stanford University measured this script in 1990 in his essay “The Dynamo and the Computer”, which went on to become famous, at a time when computers were producing exactly the same disappointment as artificial intelligence does today. To the objection of the Nobel laureate Robert Solow that one could see the computer age everywhere except in the productivity statistics, David answered with a look back: around 1900 an observer could have said the same about the dynamos. In 1899 electric motors supplied less than 5 percent of the driving power of American factories, only a good two decades later was half reached, and the measurable boost to productivity came only in the twenties, four decades after the first power station. Then, however, it came with force, because about half of the acceleration that productivity growth in American industry experienced in that decade can be traced statistically to the build-out of electric drives alone. David added a warning that strikes harder today than it did then: factory buildings wear out and eventually force a new build, but a company's processes and information structures do not wear out. Whoever does not deliberately resolve on the rebuild is forced into it by nothing.⁶⁷ Transferred to the present, that means that the real effect of artificial intelligence will not arise where employees receive one more tool, but where processes, products and responsibilities are designed from the ground up around the capabilities of digital team members.

“The first boost to productivity from electrification came only four decades after the first power station.”

“Factory buildings wear out and eventually have to be renewed, a company's processes and information structures do not. Whoever does not resolve on the rebuild is forced into it by nothing.”

From the distinction between the product and the process channel that the study by Kim and Koning contributes, a strategic decision follows that many companies have not yet recognized as one. If AI capabilities increasingly migrate into bought products, then knowledge work migrates with them out of the company and over to the vendor. Companies therefore have to decide today already, and in future far more urgently, which knowledge work they hand over to bought-in products and where they adapt the technology to their own processes so that their own knowledge stays in the company. This is not a software question but one of the most important strategic

decisions of the coming years, and it is not an either-or. The real leadership achievement consists in determining, topic by topic, which work belongs in the one channel and which in the other. That the German economy, with its pronounced inclination toward refining existing models, is already feeling its way in this direction turns the decision from a theory into an observable practice, yet it is far from being taken deliberately and thought through topic by topic.⁵¹

“If AI capabilities increasingly migrate into bought products, knowledge work migrates with them out of the company and over to the software vendor.”

What this decision looks like in everyday practice is shown by a simple pair of opposites. Payroll is a competitive advantage for hardly any company, and here the bought-in AI product is the sensible choice, because it supplies routine knowledge cheaply and reliably. Product development, by contrast, is the core of what distinguishes a company from its competitors and therefore belongs on its own side of the line. What matters there is capturing the company's expertise in a form that belongs to the company itself and that digital team members can work with directly, regardless of which model or which vendor happens to sit underneath. That way the knowledge simply travels along when the technology changes, while knowledge handed over to someone else's product stays with the vendor.

The delegation findings in this chapter show that this step is a leadership decision in its own right and not an automatic consequence of better models, and that is precisely where the good news that follows from the diagnosis lies. The bottleneck can be shaped. An organization that makes learning itself the object, that systematically examines which task it hands to a digital colleague under which level of supervision, does not shorten the long rebuild to a single quarter, but it does turn it from a fate into a program. The metric by which this report measures the distance between pilot and production year after year will show who has taken this step.

The technological foundation is being built faster than companies adopt it into their work

The three forces describe what is happening in the labor market. Beneath them runs a technological development that this report tracks all year long at workafterai.org, and this chapter condenses it into a situation report. The sources are the investment decisions of early stage investors and the work of research institutions, the same evidence base as in the annual report, because where this capital flows reveals earlier than any survey which tools will arrive in companies in the coming years. These investments say less about research than about the path of the technology into the market, because they indicate which tools are moving into everyday company work, and they therefore allow an estimate of timing. The investors are always evidence and never the subject, because what can be read from their decisions in the first half of 2026 is not about returns but about work itself, and it falls into three clearly recognizable layers.

The first layer is the transition from tool to employee. The systems into which most capital is currently flowing are no longer assistants that speed up individual steps, but digital workers that take over entire tasks from beginning to end, in customer service, in research and analysis, in insurance claims processing, increasingly specialized by function and industry. This is already visible in the size of the financing rounds, when a vendor of digital employees for companies such as Sierra (sierra.ai) receives 950 million dollars in a single round.⁸³ How fast the task length is growing that can be handed to such workers is now measured by a research line of its own. METR, an organization specialized in evaluating AI models, determines the so-called time horizon, that is, the time humans need for those tasks that a model completes with a fifty percent success rate, and this horizon has doubled roughly every seven months since 2019, in the most recent measurement period even faster.⁸⁴ The metric comes from the field of AI safety research, is limited to software and research tasks and measures under laboratory conditions rather than in everyday operations, but it is so far the cleanest translation of model progress into the language of delegable work, and this report will carry it forward in coming editions as a reference curve. Anyone who uses this curve as a reference must also disclose how fragile its upper end is. The most recent peak value of about 16 hours comes from a measurement on a preliminary version of Claude Mythos in the spring of 2026 and lies within a confidence interval ranging from 8.5 to 55 hours, because only five of the 228 tasks in the test series take that long at all. On top of that, the human comparison times come from externally paid developers, while the maintainers of the respective code base completed the same tasks 5 to 18 times faster, and the time horizon collapses outside programming, by a factor of 40 to 100 when operating graphical interfaces.²⁹ The curve therefore still shows the right thing, namely an unusually fast shift in what can be delegated at all, but it shows it for a narrow slice of work and with a yardstick that no one would accept for human work, because a success rate of 50 percent would not be a passed task in any company. What is being built here is the object around which the three forces of this report revolve, a digital workforce that companies can hire.

What happens beyond this narrow slice was examined in July 2026 by a research group around Peter Kirgis and Sayash Kapoor in the CRUX evaluation project at Princeton University, with a

method developed for the purpose. In what they call shadow evaluation, a lavishly equipped top agent receives the central research question of a high quality, as yet unpublished paper, and it has six days for it, has 3,000 dollars in model credit and its own computing capacity, can research freely on the web and commission subagents of its own; afterwards the original authors review the result the way they would review a conference submission. In both cases examined, the verdict was clear, because on the six point rating scale of a leading academic conference the machine written papers scored a two and a one. What is remarkable is where they failed and where they did not. The technical execution succeeded completely and without human help, from the literature review through setting up and repairing the computing environments to hundreds of experiments and the print ready document. The systems failed at the judgment of when a finding is sufficiently supported, at the creative response to fundamental criticism, to which they replied with further caveats in the text instead of a new experimental design, at withdrawing from dead ends once entered, at awareness of their own resources and deadlines, and at meeting the requirements, because both papers exceeded the permitted page count and would therefore have been screened out on formal grounds alone. The findings rest on two cases and on the researchers' observations of their own experiment, there is no independent replication, and two of those involved are among the best known critics of the thesis of an imminent machine self improvement, which the paper itself discloses. What it shows is nonetheless exactly the dividing line around which this report revolves, because the machine already fully masters the craft of execution, while it does not master the direction of a research project, that is, the judgment of what to do and when a result is good enough, and this finding does not come from a midsize company but from the workshop of the technology itself.⁸⁵

How far this transition has already progressed in mass market business is shown by a data set from the American housing market that Andreessen Horowitz published in July 2026 from the operations of its portfolio company EliseAI (eliseai.com), whose dialogue systems handle tenant communication for roughly one in six rental apartments in the United States. In mid 2023 the median machine handled phone call lasted 46 seconds, barely more than announcing the opening hours, at the end of 2025 about one minute and 40 seconds and thus exactly as long as a conversation with a human employee. The system that used to only take the call now conducts the conversation itself. What the same data show about the handover to the human is revealing, because if a prospect is passed to an employee immediately after the machine led preliminary conversation, a good quarter end up signing the lease, whereas if three days pass before the handover, the rate falls to under 9 percent. The value of digital labor therefore does not arise where it replaces the human, but at the seam where the two work together. The numbers are an investor's own report about its portfolio company and are to be read accordingly, but as a measurement from the ongoing operations of a mass market they show that the transition from tool to employee is already everyday practice in individual industries.⁸⁶

“The value of digital labor does not arise where it replaces the human, but at the seam where human and machine work together.”

That customer service forms the vanguard is not a peculiarity of the American housing market, it also shows up in the payment data of a great many companies. The research team of the financial

services provider Ramp, whose analysis of company data this report already draws on in the chapter on the career gap, publishes a monthly ranking of the vendors that the roughly 70,000 American companies on its payment platform are newly buying from, and in July 2026 two vendors of machine customer service systems, PolyAI (poly.ai) and Sierra, took first and fourth place on the list of risers. More remarkable than the placement itself is the reason the vendor's economist gives for it, because it reaches far beyond customer service. The success of artificial intelligence in software development, he says, can so far only be read off the number of submitted code changes, the lines written and the share of developers who work with such tools at all, and these figures show the quantity of the work but not its value. In customer service, by contrast, the success stands directly in the metrics that are tracked anyway, in the share of requests handled without human involvement and in the time to closing a case. The ranking itself is the market letter of a commercial vendor, it names no growth rates, does not disclose its selection method, and its population consists of American companies that are predominantly young and close to technology. What it nonetheless shows is a selection rule that this report encounters in several places, because digital labor prevails first where its contribution is measurable, and not where it reaches furthest technically.⁷⁶

“The success of AI in software development so far shows up only in code changes, lines written and the share of developers who work with such tools, but not in value.”

The second layer is more remarkable, because it indicates a maturing. Investors are visibly no longer asking primarily what these systems can do, but whether they can be trusted, and they are now financing a whole class of tools of a kind that so far existed only for human workforces. Separate software categories are emerging for the supervision of digital team members, for testing their reliability and for controlling what they are allowed to do, and in parallel attention is shifting to the running costs of operation, because whether an autonomous digital colleague pays off is decided not by the capabilities of the model but by the inference costs of its continuous use. Translated into the language of this report, this means that the technology industry is currently retrofitting management structure, oversight and cost accounting for the digital workforce, that is, building exactly those organizational preconditions whose absence the chapter on the long rebuild diagnosed as the real bottleneck.

A secondary measurement in the American study of AI financial advice that the chapter on the long rebuild already drew on shows how justified the question of reliability is. If the same query is put five times in a row, the recommended equity allocations vary by an average of 6.5 percentage points, while the researchers' evaluation step alone contributes only 1.6 percentage points of that. The imprecision therefore does not arise in the measurement but in the machine itself, and it is the tangible reason why oversight and reliability testing of digital workers are currently becoming a market of their own.⁸¹

“The technology industry is building the digital workforce faster than companies are learning to manage it.”

The third layer is the most recent and perhaps the most consequential, the preparation for live deployment. Digital workers are given training environments in which they are tested before they get access to live operations, instead of learning on real tasks, and at the same time the development reaches, with legal work, the first profession with genuine liability, in which an AI system takes over the volume business under the supervision of a licensed attorney and is billed not by the hour but by result. That of all fields a profession long considered barely automatable takes this step first is revealing, because the line no longer runs between simple and demanding activities, but between organizations that can regulate oversight and liability for digital work and those that cannot.

“With legal work, an AI system takes over the volume business of a profession with genuine liability for the first time.”

Two boundary conditions belong to the honesty of this situation report. The first is physical, because as soon as it is no longer individual systems but whole digital workforces that are running, infrastructure becomes the bottleneck. The investments of the large cloud providers for 2026 are estimated, depending on the definition, at 725 to 805 billion dollars,⁸⁷ the American Department of Energy expects data centers to claim between just under 7 and 12 percent of American electricity consumption by 2028,⁸⁸ and the power supply is thus developing into the hard limit of the buildout. The second boundary condition is financial, because both the Bank for International Settlements and, according to media reports, the American Treasury Department are now warning of parallels to the dotcom bubble,⁸⁹ and some of the valuations described here will not survive that test. For the question of this report, however, that is a secondary objection, because the bursting of the dotcom bubble did not undo the internet either, it only replaced its owners.

“The large cloud providers are investing between 725 and 805 billion dollars in 2026, and the American Department of Energy expects data centers to claim up to 12 percent of American electricity consumption by 2028.”

This physical limit has a second, less visible side that has less to do with building than with arithmetic. Whether a digital team member actually replaces a human colleague is decided in the end by the plain question of cost, because the only work replaced is work that is more expensive than running the artificial intelligence meant to take it over. And running it is anything but cheap. Unlike classic software, whose additional costs fall toward zero with every further user, powerful artificial intelligence incurs costs anew with every single request, for data centers, electricity, water and the data it draws on, and it therefore scales more like a streaming service that pays again for every song played than like a program that is written once and then copied a million times. As long as that remains the case, the question of replacement shifts from capability to economics, and its threshold is high at first, because the expensive frontier technology only pays off where it takes over work that is well paid anyway. The financial economist Aswath Damodaran, who initially left this cost point out of his analysis of AI disruption and expressly added it in the summer of 2026, draws the sober conclusion that the pace of the long rebuild depends not only on the learning

capacity of organizations but also on the cost structure of the technology itself, and that both brakes will only ease when the operating costs of artificial intelligence fall significantly, which is by no means settled as long as a large part of these costs remains physical in nature and hangs on electricity and computing power.³¹ That this threshold does not hang on electricity and computing power alone, but lies partly in the hands of the organization itself, is shown by the case in the chapter on the long rebuild in which the costs per unit fell while usage multiplied.

“Unlike classic software, whose additional costs per user go toward zero, artificial intelligence incurs costs for computing power anew with every single request.”

Looking ahead

Read together, the three layers show that the technology industry is currently building the supply side of the long rebuild, digital labor, its management instruments and its capacity to bear liability, and it is building faster than most organizations are developing the corresponding demand side, namely their own knowledge of which work they want to hand to this workforce and with what oversight. The bottleneck of the coming years is therefore foreseeably not on the supply side. Once the tools for trust, oversight and cost accounting are standard, the argument that the technology is not ready for live deployment loses its basis, and what remains is the question this report identifies as the real one, whether the organization is ready. Coming editions will record at this point, year after year, how far the technology has run ahead of the companies that are supposed to use it.

The organization of the future arranges its work around topics instead of functions, and it gets smaller in the process

The debate about artificial intelligence and work is almost always about which occupations are under threat, and that pushes aside the question that owners and management actually have on their desks: which blueprint a company will use to organize its work from here on. The functional organization, the division into marketing, sales, finance and engineering, is not a law of nature but a historical answer to a cost calculation. Specialist knowledge was scarce and expensive, so companies bundled the specialists into departments and accepted that work would travel through handoffs between them. Ever since the essay “The Nature of the Firm” by the later Nobel laureate Ronald Coase in 1937, we have known that the boundaries and the internal structure of firms follow coordination costs.⁹⁰ When that calculation changes, the blueprint changes sooner or later.

“The division into marketing, sales, finance and engineering is not a law of nature but a historical answer to a cost calculation. When the calculation changes, the blueprint changes.”

That calculation is exactly what is now in motion, and measurably so. The chapter on the long rebuild showed that the usage data of the three large model vendors reveal no replacement of entire occupations. What the OpenAI measurement makes visible instead is the handoff that is saved: almost half of the professional tasks the tools are used for belong historically to a different occupation, and the pattern is strongest where no specialist is within reach to delegate to.¹⁰ The first organizational effect of the technology is therefore not the replaced position but the replaced detour across a colleague's desk. Where handing the task to the specialist department used to be the norm, the task is increasingly done by the person where the need arises, with a digital team member alongside. The boundaries between roles are shifting faster than any org chart says.

“The usage data of the large model vendors reveal no replacement of entire occupations.”

“Almost half of the professional tasks the tools are used for actually belong to other occupations.”

At the level of entire companies, the same movement shows up in new ventures. The study by Kim and Koning, drawn on repeatedly over the course of this report, found that AI-native firms are set up about a quarter smaller and roughly half a hierarchy level flatter than comparable companies in the same industry and the same founding cohort, at comparable or higher valuations.⁵¹ In established companies, a broad reduction of management layers is running in parallel, which Korn

Ferry (kornferry.com) documents as a pattern without attributing it to artificial intelligence,⁹¹ and individual cases show how radical that reduction can be, for instance the reported cut in management levels in Microsoft's gaming division from 14 to three to five.²⁴ None of these observations proves on its own that the technology is the cause. Taken together, though, they suggest that organizations are moving toward fewer levels, smaller units and broader roles.

What is striking is that three very different players in the market now sketch the same target picture. From its own consulting practice, McKinsey describes units of two to five people that oversee a complete process from start to finish with 50 to 100 specialized AI agents.⁹² Gartner expects that in 2029, 60 percent of organizations will deploy small developer teams across the board, up from 15 percent in 2026, and explicitly speaks of the boundaries between roles dissolving.⁵³ Tool vendors such as Warp (warp.dev) describe the developer's new core task as understanding and adjusting the system that does the work, no longer as doing the work itself.⁹³ All three have a business interest in this picture, none of them presents a solid measurement, and their figures are expectations, not measured values. But the agreement between consulting, market research and software vendors with what the usage data already show about employee behavior means the target picture has to be taken seriously: the small unit built around a topic, fully accountable for an outcome, pulling the adjacent expertise from the machine.

“The target picture is the small unit built around a topic, fully accountable for an outcome, pulling the adjacent expertise from the machine.”

What it looks like when a corporation applies this target picture to itself has been on the record since August 2026 at Meta. The restructuring plan that the first chapter treats as the fourth case of the AI argument was, at its core, an organizational plan. According to the guide circulated internally, small units of two to three developers and one designer were to work in four-week cycles instead of half-year plans, the job titles product designer and software developer were to be merged into the single title builder, the middle management layers were to disappear, and every unit was to report to a single area lead who decides on evaluation and promotion for 30 to 50 people, while the head of the small unit no longer had any management authority; an agent-supported analysis was to set the daily priorities. By June, according to the Reuters report, at least eleven areas had adopted this form. That is, almost point for point, the target picture drawn by consulting, market research and tool makers, implemented in a corporation with 75,000 employees. What Reuters quotes from the internal forums is the cost side that this chapter describes further below: an employee who had been given the lead of a unit noted that he received neither management training nor access to the evaluation tools; the responsibility had been distributed, the authority had not. Together with the internally measured incidents, which rose by 40 percent, and a workforce whose approval fell from 74 to 55 percent within half a year, that adds up to the first documented large-scale attempt to build this chapter's target picture, and its provisional outcome: the downsizing could be ordered, the coordination could not be abolished.²⁶

What it looks like when this target picture meets a single function can already be inspected in marketing, the function whose tasks travel furthest into other occupations according to the OpenAI measurement. A McKinsey survey of 521 marketing leaders published in June 2026, produced in

cooperation with Google's advertising business and therefore to be read with a double interest at play, traces there the same movement this chapter describes. Marie Gulin-Merle, responsible for global ads marketing at Google, sums it up in the sentence that AI agents do not respect silos, and expects that responsibility for performance marketing, growth and brand will no longer sit in separate roles. The consultancy itself sketches the future marketing organization as a diamond rather than a pyramid, explicitly as an illustration from client examples and not as a measurement, with fewer purely executing people and three role types that repeat this chapter's blueprint almost word for word: some build the digital team members, others orchestrate them across the entire process, and the third group secures standard and judgment as the quality authority. More telling than the target picture is the measured distance from it. Only 28 percent of respondents see a fundamental rebuild of teams and processes in their own company, and while just under 60 percent use the tools several times a week, fewer than 10 percent draw a benefit from end-to-end processes, which means the gap between pilot and production from the chapter on the long rebuild returns here in marketing. One figure, finally, exposes the emotional side of the organizational decision: the tension is greatest at the very top, with 96 percent of marketing chiefs reporting high enthusiasm, 71 percent reporting fear at the same time, and 80 percent seeing a risk to their own role. The reordering of work is therefore decided by the people most exposed to it.⁹⁴

“Marketing tasks travel furthest into other occupations.”

“Marketing chiefs report high enthusiasm and, at the same time, fear for their own role. The reordering of work is decided by the people most exposed to it.”

That this movement does not stop at the company boundary becomes clear when you look at the supply side of marketing. Martin Sorrell, who built WPP (wpp.com) into what was for a time the largest advertising group in the world and who today runs the digital agency group S4 Capital (s4capital.com), told his own industry in late July 2026 that it is having the wrong discussion. The biggest disruption does not lie in the question of whether the machine replaces the creatives; the real change runs through media, data, platforms, organizational structures and the economics of the agency business. Clients will demand that the efficiency gains of the technology be passed on to them, agencies that continue to sell time rather than outcomes will come under growing pressure, and the large platforms will become even more powerful, while retail advertising inventory and the influencer business shift where the ad money flows. This is not a measurement but the assessment of a man who earns his living with S4 Capital on the counter model and who presents it as a pointer to his own podcast appearance. It is revealing nonetheless, because it describes from the supply side the same pincer movement that the usage data show from the demand side. From the inside, neighboring functions take over the tasks of the function that, according to the OpenAI measurement, gives away its work furthest, as when sales builds its own sales campaign with digital team members and management pulls brand stewardship to itself. From the outside, the platforms and the retailers take over the execution that used to sit with the agencies, whose business model of selling working time bills for precisely what the technology devalues fastest.¹³

This outer flank can now be quantified. The advertising channel that structurally bypasses the agency, because the retailer markets its own inventory directly, grew to just under 175 billion dollars in 2025 according to the forecast of the market researcher WARC and will, at just under 197 billion in 2026, overtake combined spending on classic and connected television for the first time; WPP's own media holding GroupM arrives independently at almost the same figure, and the fact that of all players the largest agency holding documents the growth of the channel that undercuts its intermediary role speaks for the number. Depending on the count, that puts roughly 15 to 16 percent of the entire advertising market in retail advertising inventory. At the same time, according to the same source, Alphabet (abc.xyz), Amazon (amazon.com) and Meta together account for around 55 percent of the advertising market outside China, and the trend is rising; Amazon's advertising business grew by 26 percent in the second quarter of 2026 to 19.8 billion dollars, Meta's by 27 percent to 59.4 billion. Where the platforms are taking this business was described unmistakably by Mark Zuckerberg in an interview in May 2025: a company tells the platform its goal and provides its bank details, and the system handles creation, audience selection and measurement. Google has already taken the first step in that direction and has since August 2026 been automatically moving its clients' search campaigns into the AI-driven campaign format. How far reality is keeping up is shown by an April 2026 survey of media agencies by the trade service Marketing Brew: automated delivery already accounts for 60 to 70 percent of Meta budgets at individual agencies, yet none of the agencies surveyed knew of a client that already fully trusts machine creation. The gap between pilot and production of the long rebuild returns on the platform side.¹⁴

How this double pressure affects the supplier industry can be read in its own financial results, and those show no simple decline but a spread. WPP, for decades the industry benchmark, lost 5.4 percent organically in 2025, closed the year with a net loss, cut the dividend from 39.4 to 15.0 pence and, after a profit warning in July 2025, lost 16 percent of its market value in a single day, to its lowest level since 2009; the new chief executive Cindy Rose, brought in from Microsoft, has been rebuilding the group since. Publicis (publicisgroupe.com), by contrast, which bet early on data and its own AI platform, grew 5.6 percent organically in the same year and raised its prediction in the summer of 2026. The merged group of Omnicom (omnicomgroup.com) and Interpublic (interpublic.com), the largest vendor in the industry since November 2025, is growing and cutting staff on a large scale at the same time; revenue and employment are therefore decoupling even where the business is going well. And Sorrell's own group S4 Capital, launched as the digital counter model, is shrinking for the third year in a row and owes its profit jump in the first half of 2026 almost entirely to cost cutting. The crisis is therefore not hitting the agency as a species but the business model of selling working time, and it spares not even the person who diagnoses it most loudly.¹⁵

On the client side, the pincer closes. The American association of advertisers ANA has been measuring for almost two decades how many of its members run an in-house agency: 42 percent in 2008, 58 percent in 2013, 78 percent in 2018, 82 percent in 2023, with 92 percent still working with external agencies at the same time, so the shift concerns tasks, not the relationship as such. What is new is that artificial intelligence is accelerating this shift. In the spring 2025 spending survey by the market researcher Gartner, 22 percent of marketing chiefs state that generative AI has already reduced their dependence on external agencies for creation and strategy, and 39 percent plan to cut their agency budgets. What that looks like in individual cases is demonstrated

by two companies whose figures are to be read as self-reported: the payment provider Klarna (klarna.com) reported as early as spring 2024 that it had cut its spending on external agencies by a quarter and shortened the production cycle for campaign images from six weeks to seven days, and Unilever (unilever.com) has been building its own AI design studios since the summer of 2025, remarkably together with the agency group Interpublic, which thereby earns from the move of value creation to its own client. The academic survey of marketing chiefs by Duke University adds the staffing side to the picture: growth in marketing teams has halved within a year, and respondents expect the technology to take over more than half of marketing activities within three years. Here too the counter calculation belongs in the picture: the British industry survey by ISBA measured a considerable gap as early as 2023 between what companies expected from the shift into their own teams and what they experienced. In-house it is often no faster, the paths are merely shorter.¹⁶

Taken together, that is the full pincer: from the inside, neighboring functions take over marketing tasks with their digital team members, from the outside, platforms and retailers pull the execution to themselves, and on the client side, whatever agency work remains moves into in-house teams. What is left for the function in the middle is the role the McKinsey typology describes as the quality authority, and with it the question of whether that still requires a department of its own. Of all the classic functions, marketing is therefore the first in which this chapter's blueprint can be inspected in full, from the inside as well as from the outside, and therefore also the one in which it will be decided earliest whether a function shapes the rebuild or suffers it.

Anyone who turns this into a simple downsizing formula, however, is leaving out the cost side of the rebuild. Coordination does not disappear, it shifts, from alignment between people to oversight of parallel digital contributors, and there are precedents for this in the history of failed experiments with flat hierarchies, from Zappos (zappos.com) to Medium (medium.com), in which it turned out that the abolished coordination had not been saved but merely shifted in silence.⁹¹ The orchestration load of the previous chapter is the human side of this calculation. The business side has been measured quarterly by KPMG: 49 percent of the organizations surveyed have already adjusted the pace of their AI rollouts once expected costs began to exceed expected benefits, and only a good third have full visibility into their own AI operating costs.⁷² On top of that comes a quieter danger, which the Boston Consulting Group describes as a distributed loss of expertise: when everyone does a bit of everything, the depth of expertise on which the quality of the whole rests can erode.⁵⁶

The OpenAI measurement already contains this second warning itself, because what is measured are attempts, not results. Whether work done outside one's own field delivers what it promises is not shown by usage data, and OpenAI explicitly notes that specialists remain central for judgment and review.¹⁰ An organization that simply lets its employees cross boundaries without reordering review and accountability is not organizing productivity for itself but an unanticipated risk. The Stanford research on human-centered AI supplies the framework for this: the division of labor between human and machine has to be designed according to task, expertise and risk, and the risks accumulate over time instead of becoming visible in individual errors.⁹⁵

“An organization that simply lets boundary crossings happen is not organizing productivity for itself but an unanticipated risk.”

From both follows the real organizational decision of the coming years, and it is a decision no one is spared. The delegation experiment by Colaiacovo and Koning showed that the willingness to hand work to the machine depends only to a good third on what one credits the technology with, and beyond that is an independent leadership decision.⁷ The same holds for the structure of the organization. Three sub-questions belong to this one decision: first, which value creation a small team will be accountable for from start to finish, with digital team members for the adjacent competencies; second, where depth of expertise deliberately stays bundled, no longer as an executing department but as an authority for standards, review and judgment; third, from the chapter on the long rebuild, which knowledge work may migrate into the purchased product and which stays tied to the company's own processes. Whoever does not make this decision gets the new division of work anyway, just without steering it, because employees are already crossing the old boundaries today, chat by chat.

Looking ahead

If you extend the lines of this chapter, what emerges is the picture of an organizational type that differs from the org chart of the industrial age more clearly than anything digitization has produced so far, with small units arranged by topics and outcomes, each leading people and digital team members together, while the classic functions become what quality authorities always were, guardians of standard, review and training. Careers in this organization lead less often through the size of a span of control and more often through the depth of the topics one is accountable for. Whether it turns out this way is open, and the history of management fads counsels caution. This assessment will therefore track the question of the division of work annually from now on, from 2027 in the board pulse and, from this edition onward, in its own predictions as well. The companies that shape this division deliberately will differ from those that let it happen, and that is in the end the same fault line along which this report pins the future of work as a whole: it is not the capability of the machine that decides, but the capability of the organization to reorganize itself around it.

“The organization of the future will differ from the org chart of the industrial age more clearly than anything digitization has produced so far.”

“Careers in these organizations lead less often through the size of a span of control than through the depth of the topics one is accountable for.”

Anyone who transfers the American debate to Europe is measuring with the wrong yardstick

Now that the three forces have been told apart, a question presses forward that the evidence so far does not answer at all, because that evidence consists almost entirely of American data. Does the same picture hold for Europe, or is an American story being quietly transferred to a continent with different rules? The obvious assumption, backed by the force of the Stanford numbers, would be that the same collapse at the entry level must show up in Europe as well, as soon as one looks closely enough.⁴ The European aggregate figures, however, speak against that assumption at first glance.

At the level of the 27 EU member states, the youth unemployment rate stood at around 15 percent in April 2026, the overall rate at 6 percent. Over the course of the years since 2022, overall unemployment has fallen slightly and has held stable at a historic low, while youth unemployment has risen from just under 15 to a good 15 percent, so the relative disadvantage of the younger group has widened slightly, even though the labor market overall remains in good shape. But a different number is the decisive one. The American finding of a collapse among young university graduates does not appear in the Union's aggregate figures; on the contrary, the employment rate of younger graduates rose to around 83 percent in 2025, after a little more than 82 percent the year before, with a small dip among university graduates that does not continue. Anyone who transfers the American headline to Europe unchecked is therefore measuring with a yardstick that the European numbers themselves do not supply.⁹⁶

That does not mean nothing is happening in Europe, but that the average of the whole continent hides what matters. Germany, for instance, has the lowest youth unemployment rate in the entire European Union at around 7 percent, well below the EU average, while Austria stands at around 11.5 and Switzerland at 8 percent.⁹⁶ The large European institutions, Eurostat, the OECD and the European vocational training center Cedefop, have so far not confirmed any collapse at the entry level along American lines, and the fact that it has failed to appear is itself a finding and not merely a lack of evidence.^{97, 36, 98}

An honest reading of this non-finding, however, has to include a second possibility that qualifies the all clear. The absence of the American pattern can mean protection, but it can also mean delay, because the technology is arriving in Europe later and flatter than in the United States. The continent's private AI investment stands at a fraction of the American figure⁹⁹, the number of first time AI funding rounds for young companies has grown eightfold within four years in the American data⁵¹, while the European founding landscape remains distinctly narrower, and even where European companies have long been using the tools, the actual intensity of use is growing only slowly according to the Bundesbank data.⁶⁶ But if the American entry level collapse, whichever of its three possible causes lies behind it, is connected to the diffusion of the technology, then today's European aggregate figures may be measuring not the resilience of the continent but only its lag in diffusion. European regional research supports this reading with a look below the country level. The

regional economists Andrés Rodríguez-Pose and Federico Bartalucci, in their paper “A Union under siege, or a Union besieging itself?” published in June 2026, set the AI exposure of Europe's regions against their innovation capacity and come to the conclusion that large parts of southern and eastern Europe appear untouched by the technology only because they are largely excluded from the knowledge economy. Apparent safety from automation, the authors write, is not resilience but irrelevance.¹⁰⁰ The European finding would then be no counterproof but a window of time, and whether it closes is among the questions this report will keep following in the same numbers in the coming editions.

“Apparent safety from automation is not resilience but irrelevance.”

What diffusion looks like in a single European country can be demonstrated with Austria since the summer of 2026, because the AI-Econ Lab at Örebro University, whose exposure index plays a role of its own in the next chapter, now publishes its indicators.¹⁰¹ Three mutually independent surveys produce a picture that does not fit together on its own. The share of employees whose occupation falls in the most exposed quarter of all occupations is 32 percent in Austria, just above the mean of 30.2 percent across 36 countries, but lower than in Germany at 35 and Switzerland at 36 percent. The share of companies actually using artificial intelligence stands at 30 percent according to the Eurostat survey for 2025, well above the European Union average of 19.9 percent and above the German figure of 26 percent. The share of job postings that ask for AI skills, by contrast, is 0.8 percent, third from last among 22 measured countries against a mean of 1.9 percent. Austrian companies therefore use the technology more often than average and still post almost no jobs that ask for skills in it. The three numbers count different things, namely employees, companies and postings, and cannot be set against one another; taken together, though, they describe exactly the situation this report calls the long rebuild. The tools have been bought, the work has not yet been reorganized, and as long as it has not been, the technology shows up neither in the job profiles nor in the productivity figures.

“AI skills are asked for in 0.8 percent of Austrian job postings, third from last among 22 measured countries.”

The measurement of Austria's specialist AI workforce by the team at the Vienna University of Economics and Business, already introduced in the chapter on the career gap, adds a fourth measurement to this triad that explains why it is so hard to resolve.⁵⁰ Counted at the level of individual employees, Austria's specialist AI base is both unusually high in quality and unusually narrow. Measured by the share of specialized roles among all AI roles, the country ranks third out of 38 European countries, behind Germany and Switzerland; measured by density, that is, AI specialists per 1,000 employees, only 19th. At the same time, 97 percent of the companies observed in the panel employ nobody at all in an AI role, and even in the size class of 11 to 100 employees only a good one in eight companies has someone in such a role. That gives the contradiction in the triad an explanation: a country whose companies use the technology more often than average but employ almost no staff of their own to carry it is using it as a bought in tool, not as a capability of its own, and consequently has little reason to post jobs for it either. The numbers come from the same panel of profiles from professional networks and are to be read

accordingly as values of that panel; they do not merely confirm the reading of the long rebuild, they add to it the observation that the rebuild cannot even begin without people of one's own.

How far this reading has already worked its way into official prediction arithmetic is shown by a look at the medium term prediction issued by the Austrian Institute for Advanced Studies in July 2026. In its section on America, it names productivity gains from the use of artificial intelligence as the most important future growth driver of the United States, with cited estimates of between half a percentage point and one and a half percentage points of additional annual productivity growth, factored into expected growth of 1.8 percent per year. In the prediction for its own country, by contrast, the same technology does not appear in the baseline scenario at all, but shows up only in the risk section as an opportunity that could improve the outlook, while Austria's trend productivity for the coming five years is set at 0.1 percent per year, after the country's total factor productivity has not grown at all since 2023. Whether the American calculation is too optimistic or the European one too cautious will only become clear later; what is telling is the asymmetry itself, because the same economists credit the same technology with a measurable contribution to growth on one continent and with none as yet on their own doorstep. The lag in diffusion is therefore no longer only in the investment figures, but in the growth arithmetic of the predictions themselves.¹⁰²

“The same institute names artificial intelligence as the most important future growth driver for the United States; in the prediction for its own country it appears only in the risk section as an opportunity.”

The real European controversy therefore has to be sought where it is actually being fought out, namely in the individual DACH countries themselves.

The real European controversy is being fought out in the individual countries, not in the aggregate figures of the union

On German soil at least part of the American pattern is taking shape, even if the attribution of causes remains as contested as in the United States itself. The German unemployment rate among graduates climbed from just under 3 percent in 2024 to just over 3 percent in 2025, crossing that threshold for the first time since 2007, with around 335,000 unemployed graduates and an increase of 16 percent over the previous year. The Federal Employment Agency itself, however, names not artificial intelligence but economic weakness and structural change as the cause, and a good fifth of the increase is due to refugee migration, while graduates continue to hold the lowest unemployment rate of all qualification groups.⁹⁶ The job board StepStone (stepstone.de) reports for the first quarter of 2025 a share of entry level positions that is 45 percent below the five year average, also-called junior paradox, yet this figure too is overlaid by the weak economy and cannot be cleanly attributed to the technology.¹⁰³

Since early August 2026 this German picture can also be drawn occupation by occupation. The Ingenieurmonitor, compiled by the Verein Deutscher Ingenieure (VDI) together with the Institut der deutschen Wirtschaft, counts a good 58,000 unemployed people in engineering and IT occupations for the fourth quarter of 2025, the highest level since the survey began in 2011 and an increase of 16.7 percent over the previous year, while the number of open positions fell by 18.2 percent to just under 97,000.¹⁰⁴ What is revealing is how unevenly this decline is distributed, because it follows the same pattern as in the United States: in computer science occupations a third of the postings disappeared, in civil engineering occupations not even 4 percent, and supply shrank most sharply in Bavaria and Baden-Württemberg, where the ailing vehicle industry shapes demand.¹⁰⁴ The Federal Employment Agency reports the same picture for IT occupations, the stock of registered vacancies fell by 22 percent in 2025 to around 13,000, and the occupation specific unemployment rate rose from 3.7 to 4.5 percent.¹⁰⁵ Both institutions nevertheless read their own figures explicitly as a consequence of the business cycle and not of the technology. The VDI warns against deriving a falling need for technical specialists from the slump, and the Federal Employment Agency points out that employment continues to grow, to 1.15 million ICT specialists in jobs subject to social insurance contributions in 2025, 2 percent more than the year before. What its executive board chair Andrea Nahles describes alongside this, however, is not a mere cyclical trough but structural change made up of rising skill requirements, more academic qualification and more specialization, in which experts with an unemployment rate of 3.5 percent are barely affected.¹⁰⁵ On this reading demand is not disappearing, it is shifting toward the specialists.

How German companies themselves answer the questions that were also put to American ones has been shown for the first time since the spring of 2026 by a survey involving several central banks. For the working paper "Firm Data on AI" by Yotzov, Barrero, Bloom, Bunn, Davis and others, with which the first chapter already documented the expectation gap between leadership and

workforce, the Deutsche Bundesbank put the same questions in the same wording to firms in its online panel in January 2026 as the Federal Reserve Bank of Atlanta, the Bank of England and an Australian university did in their panels. Around 2,500 companies responded. The German picture is subdued throughout. 65 percent of German firms use at least one AI technology, against 78 percent in the United States. 95 percent report no effect on their employment over the past three years, against 89 percent there, and for the coming three years German firms expect a considerably smaller change in their employment than American ones, which on average reckon with minus 1.2 percent, something the authors cautiously connect to the more heavily regulated German labor market. Up to this point the measurement supports the argument of the following chapter, according to which protection against dismissal makes the same shock arrive more quietly in Europe. The same damping, however, has an uncomfortable flip side. On productivity, German firms expect growth of 0.9 percent over the coming three years, American firms 2.3 percent and British firms 1.9 percent, and 91 percent of German companies have so far observed no productivity effect at all. What appears as protection on the employment side appears as a lag on the productivity side, and precisely this lag touches on the converse of the demographics chapter, according to which the technology can offset the departure of the baby boom cohorts if bold investment is made now. The report lets this tension stand instead of smoothing it over. The expectation gap between 0.9 and 2.3 percent is perhaps the single most precise figure for the delay hypothesis of the previous chapter, and it says nothing about whether the lag can be made up, only that the companies themselves are already reckoning with it.⁵⁴

The methodologically most convincing and at the same time most European evidence comes from Switzerland. The KOF study by ETH Zurich explicitly replicates the Stanford method and shows that the number of registered job seekers in occupations highly exposed to AI has risen by up to 27 percent more than in occupations with low exposure, an effect that applies particularly to younger job seekers. For those under 50 in highly exposed occupations the number of unemployed has even more than doubled since the autumn of 2022.¹⁰⁶ At first glance this is contradicted by the European Central Bank, whose analyses of its own SAFE business survey from two survey waves in 2025 show no significant difference in the hiring balance between firms that use AI and those that do not, with a slight tendency even toward hiring rather than dismissing.¹⁰⁷ Both investigations are methodologically sound, and both simply measure at different levels, the Swiss study at the occupations of job seekers, the central bank at the hiring balance of firms across the entire euro area.

“In occupations highly exposed to AI the number of unemployed people in Switzerland grew up to 27 percent faster than in occupations with little exposure, and among those under 50 it has more than doubled since the autumn of 2022.”

How quickly this controversy falls back into the confusion of exposure with replacement, against which the studies on the career gap warn, could likewise be observed in Switzerland in early August 2026. The Tages-Anzeiger laid the Dynamic AI Occupational Exposure Index, an annually updated exposure measure from a research team around the AI-Econ Lab at Örebro University, which also includes the Kiel economist Holger Görg,¹⁰⁸ over 391 occupations in the Swiss labor

market and asked in its headline how replaceable one's own occupation is by artificial intelligence and which occupations are most at risk.¹⁰⁹ That is precisely what the index, on the explicit statement of its authors, does not answer, because it captures the potential applicability of the technology to the content of an occupation and leaves open whether displacement, complementarity or nothing at all follows from it, since that depends on how widely the tools spread, on complementary skills and on the decisions of organizations. The case becomes revealing because the same team used register data covering almost all firms in Denmark, Portugal and Sweden over two decades to test what consequences exposure, measured with its own index, has for employment. According to the authors' summary this analysis found no systematic effect on the total employment of more heavily exposed firms, but it did find a consistent upskilling, because demand shifts from simple clerical work to highly qualified white collar work.¹⁰⁸ For the European controversy this is revealing in two respects. Alongside the central bank's hiring balance, there is now a second firm based European measurement that shows no volume effect but a structural effect, and the Swiss case at the same time demonstrates how a measure of possibilities turns into a headline about risk. The article itself does concede that a high index value does not mean a large risk of dismissal, but what remains is the ranking of risk, not the qualification.

Together with the Austrian finding already mentioned, according to which only a very small share of the open positions advertised as data or AI roles goes to people at the entry level⁴⁷, what emerges is indeed a distinct version of the American dispute between Stanford and the Federal Reserve, fought out on European soil, and not a mere copy of it. It is most clearly visible in Switzerland, in Germany more hinted at than proven, and it has still not reached the official EU statistics.

Dismissal protection is the structural reason why the same shock lands more quietly in Europe

Why the same technology lands markedly more quietly in Europe than in the American wave of layoffs described in the first chapter is explained not by the technology itself, but by a structural feature of the continent: dismissal protection. The relevant OECD index does exist in several, in part contradictory versions, yet the basic picture is stable across all of them. In the classic version, the United States, with an index value close to zero, comes last among all 42 countries covered, while Germany stands at about 2.5, Austria at about 2.25 and Switzerland at about 1.5, against an OECD average of a little over two. In almost every American state an employment relationship can be ended at any time and without statutory severance, whereas in Germany the Dismissal Protection Act applies after half a year of service in companies with more than ten employees, with graduated notice periods and a de facto severance standard, and in Austria there are graduated notice periods as well as the continuously accrued new severance system.¹¹⁰

The most convincing empirical evidence for the different European adjustment path is short-time work. In Germany it reached its peak in April 2020 with about 6 million employees, in Austria about one million people, just under 30 percent of everyone actively employed in the country.¹¹¹ The result was that Europe kept unemployment low in one of the sharpest economic crises of recent history, while in the United States it briefly shot up into double digits. The same logic applies to any technological shock that an AI could set off. Where American firms lay people off quickly and hire again just as quickly, Europe holds on to staff through short-time work and is pushed harder by law toward retraining than toward separation, which is consistent with an older McKinsey survey from early 2018, to be cited with caution, according to which 94 percent of the European executives surveyed, against 62 percent of the American ones, rely on retraining rather than on new hiring, a figure from the time before generative AI that nonetheless shows the same structural pattern.¹¹² What this path looks like inside a company is shown by Europe's largest furniture retailer. That the same group cut about 800 office jobs in March 2026 and that the franchisor Inter IKEA (inter.ikea.com) cut a further 850 in May, in both cases without any reference to the technology, belongs to the same picture, because what was protected was the customer side and what was thinned out was the administrative level.¹¹³ Exactly this difference is the reason why the wave of layoffs described in the first chapter has so far run in a more muted way in the DACH region than in the United States, not because European companies were more restrained, but because quick dismissal is made harder for them legally and financially.

What is remarkable is that this reading is by now being confirmed from outside, and from the direction one would least expect it. The American entrepreneur Dan Shapiro, head of the laser device maker Glowforge, reduced the decision facing companies with more productive employees to three options at the end of July 2026. The first is to do nothing, the second to pursue the same targets with half the workforce and pocket the savings, and the third to keep the workforce and double the output instead. And he expressly records that countries with strict dismissal protection do not have the second option at all, which is why only the third is left to them. What counts in the

European debate as a disadvantage for the location thus appears from an American point of view as a forced advantage, and the question Shapiro attaches to it is the genuinely interesting one: who wins when a halved company competes against a doubled one?¹¹⁴

The thought should not be confused with its premise, though. Shapiro reckons with a doubling of individual effectiveness, and for that there is no evidence worthy of the name. The two field experiments that the chapter on the long rebuild cites by name arrive at a quarter to two fifths in time savings on individual, clearly defined tasks, and in the same trial they show that the success rate outside that range falls by 19 percentage points instead of rising.^{9, 61} What remains is not the number but the switch. In Europe it is set differently than in the United States, and what set it there was not insight but the law.

Demographic pressure inverts the whole starting question

While the American debate assumes a labor market that has too many workers, workers a machine could make redundant, Europe faces the opposite starting position, and this inversion is the real reason why the job killer story stands on weaker legs here than elsewhere. According to the most recent release from the Federal Statistical Office in June 2026, Germany will lose around 13.3 million people in the labor force to the statutory retirement age of 67 by 2040, that is 30 percent of all people in the labor force who were available to the labor market in the year before, and the younger cohorts cannot make up this loss in sheer numbers. The two oldest cohort groups alone, those shortly before retirement, comprise around 10 million people in the labor force, more than the entire group of 25 to 34 year olds, and the share of those over 55 in the total labor force has risen from a little over a fifth in 2015 to more than a quarter today.¹⁷ The old-age dependency ratio has been rising for years across the entire European Union, but the effect is felt most where, as in Germany, a single, exceptionally large generation leaves as a block within a few years. That this decline is not a distant projection but has already begun is recorded by the Institute for Employment Research of the Federal Employment Agency in its labor market prediction, most recently confirmed in spring 2026: in 2026 Germany's labor force potential falls for the first time ever, by around 40,000 people, because the exit of the large birth cohorts is for the first time no longer offset as immigration weakens.¹¹⁵

“By 2040 Germany will lose 30 percent of its labor force to retirement, and the younger cohorts cannot make up this loss in numbers.”

Austria shows the same picture on a smaller scale. According to a December 2025 study by the consultancy Prognos (prognos.com), the number of people of working age falls from around 5.5 million in 2023 to around 5.3 million in 2035, and if the economy keeps growing over this period as expected, around 246,000 workers will be missing by then, an order of magnitude also supported by a study by the Vienna Institute for International Economic Studies commissioned by the Ministry of Economic Affairs, according to which around 540,000 workers will leave working life through the retirement of the large birth cohorts between 2022 and 2027 alone.^{116, 117} How directly this mechanism already shapes the current decade is shown by the medium-term prediction of the Institute for Advanced Studies from July 2026. It expects the Austrian unemployment rate to fall from 7.5 percent this year to 6.4 percent by 2030, even though the economy grows by only one percent a year on average over the same period, because the supply of labor barely rises any more for demographic reasons and the working-age population, according to the underlying projection by Statistics Austria, grows at all only through immigration; without immigration it would shrink by 2030. The shortages of skilled workers, the institute states explicitly, will intensify further. In this prediction the Austrian labor market therefore does not ease because enough new jobs are being created, but because the people who could fill them are becoming scarce.¹⁰² Switzerland too is aging in every population scenario of the Federal Statistical Office, and in the more cautious of

these scenarios the domestic working population could no longer be sufficient from the early 2050s onward to cover the economy's demand. An analysis by the umbrella business association *economiesuisse*, to be classified accordingly as the voice of an association, puts the gap at around 431,000 missing workers by 2040, that is around 8 percent of those employed today, with the peak of the retirement wave in 2029, when for the first time more than 30,000 more people reach retirement age than young ones move up.¹¹⁸ Since spring 2026 the federal government itself has stood behind this diagnosis, because a report based on a study by the State Secretariat for Economic Affairs and adopted by the Federal Council in April names demographic aging as the main cause of the Swiss labor shortage.¹¹⁹ For comparison, in the United States the population keeps growing well into the 2030s according to the projections of its own statistical agency, because immigration there is increasingly becoming the load-bearing engine of growth and keeps the share of people of working age stable, while Germany, Austria and Switzerland are structurally heading toward a shrinking working population without comparable immigration reliably offsetting it.¹²⁰ In the speech already mentioned, Bundesbank President Joachim Nagel named exactly this possibility of cushioning demographic pressure, because AI-driven productivity gains could offset or at least soften the losses from the decline in the working-age population.⁶⁶ The starting question of the job killer story therefore poses itself exactly the other way around, because the real problem of an economy that loses almost a third of today's labor force over the coming decade and a half is too few workers, not too many.

• THE DEMOGRAPHIC PRESSURE

The exit of the baby boom cohorts hits all three countries.

Three measures, not comparable with each other

Germany	Austria	Switzerland
13.3M	246,000	431,000
people in the labor force reach retirement age by 2040, 30 percent	workers missing by 2035	workers missing by 2040, around 8 percent
Destatis, June 2026	Prognos, December 2025	economiesuisse, industry group source

Source: Statistisches Bundesamt (June 2026), Prognos (Dec. 2025), economiesuisse (industry group).

FIGURE 6

“While Germany, Austria and Switzerland are heading toward a shrinking working population, the American one keeps growing well into the 2030s.”

These country figures still say nothing about who the departure hits first, and that is exactly where the picture becomes interesting, because it contradicts the common assumption. The official age structures of the three countries line up into a remarkably stable ranking. At the top everywhere is public administration, in Germany with 29 percent of employees over 55 against 23 percent in the economy as a whole, in Austria with 24 percent and, with nearly 150,000 people affected, at the same time the largest absolute stock of any sector, in Switzerland with 22 against 19 percent. Behind it come transport and storage as well as utilities and waste management. At the lower end stands the same group in all three countries, namely information and communication, hospitality and business services, and the distance is considerable, because Austrian information technology services come to 10 percent and thus to less than half the top value. Contrary to expectation, health and social services is not at the top but, in Germany, exactly at the average, because its gap arises not from departures but from the simultaneous additional demand, which in the predictions of both countries accounts for by far the largest increase in employment.^{121, 122, 123}

“At the top everywhere is public administration, it loses staff to age-related departures first and most.”

Most revealing is construction, because it shows that the sector level is too coarse. The construction industry is not above average in age in any of the three countries, in Germany it is below the average at 21 percent, in Switzerland likewise at 18 percent, and in Austria civil engineering stands almost at the end of the entire list among those 60 and older. Go one level deeper, to occupations, and the picture flips completely. In building construction supervision, 41 percent of employees are 55 or older, in civil engineering 37, and there, according to the same analysis, nine out of ten open positions can no longer be filled. Logistics shows the same pattern, where the driving occupations are at 40 percent and above, while the sector as a whole is only moderately above the average. It is not the sectors that are aging, but their management and supervisory levels. From this follows a feature of the European picture that has so far received the least attention. The demographic departure hits first and hardest the work that depends on experience, supervision and local knowledge, that is administration, construction site management, driving, maintenance as well as utilities and waste management, and it hits last the work where artificial intelligence is put to use first. Where the people disappear, the technology reaches least far, and where it reaches furthest, the fewest people disappear.^{121, 122, 123}

“It is not the sectors that are aging, but above all their management and supervisory levels.”

That this ranking also shapes companies' wage expectations is shown from the other side by a survey by the ifo Institute from June 2026, because construction companies are the least likely to expect that the use of artificial intelligence will push wages down, at 39.0 percent for short work experience and 31.3 percent for long work experience, while among service providers the figures are 53.3 and 44.2 percent, and conversely it is precisely the construction firms that most often expect rising wages for their qualified staff. Where experience becomes scarce, companies trust the technology least to make it cheaper.⁴⁹

Looking ahead

It is precisely here that artificial intelligence could turn from the supposed problem into the answer. If the DACH region loses millions of workers over the coming decade and a half, then a technology that raises productivity per head and takes over routine work is not a threat scenario but the most likely instrument for maintaining the prosperity of a shrinking working population, exactly the cushioning effect the Bundesbank describes.⁶⁶

“If the DACH region loses millions of workers, the technology is not a threat scenario but the most likely instrument for raising productivity per head and maintaining prosperity despite a shrinking working population.”

This hope is of course tied to a condition, and it is not met as long as the productivity effect stays as small as the most cautious calculations show it to be. In his macro analysis, the economist Daron Acemoglu arrives at an increase in total factor productivity of 0.53 to 0.66 percent over ten years, that is around 0.06 percent a year, and the budget model of the University of Pennsylvania puts the contribution of artificial intelligence to American productivity growth for 2025 at 0.01 percentage points.²⁹ Magnitudes like these do not offset the loss of a substantial part of the labor force. The European calculation therefore only works out if the technology reaches the upper end of the published estimates, and it is thus less a prediction than a task.

The window for this, however, is not a comfortable one. The large birth cohorts that leave by 2040 carry a large part of the experiential knowledge of these economies, and only while they are still in the business can this knowledge pass into processes, systems and digital team members instead of leaving with them. Whoever invests boldly while the experienced are still there turns demographic pressure into the strongest reason to invest in this technology that exists in Europe. Whoever waits until the gap is there will no longer fill it with any machine, because by then no one will be left who can teach it what to do.

“The large birth cohorts that leave by 2040 carry a large part of the experiential knowledge of these economies, and only while they are still in the business can this knowledge pass on.”

European sovereignty is a strategic lever, not a compliance argument

From the same speech by Nagel, the keynote “Already Here, Not Yet Everywhere” of April 2026, follows the last and least discussed building block of the European picture, namely that the much invoked European sovereignty in artificial intelligence is not a compliance argument but a genuine strategic lever. For the year 2024, the Stanford AI Index counts 40 significant American AI models against 15 Chinese ones and merely three European ones, and the capital gap behind that has widened further since. According to the 2026 edition of the same index, private AI investment in the United States reached around 286 billion dollars in 2025, against around 12 billion in China and around 21 billion in all of Europe. These figures capture only private equity capital in AI firms, that is, venture capital and stakes, and not yet the infrastructure spending of the large cloud groups, which at an estimated 725 to 805 billion dollars for 2026 alone is in a category of its own and is borne overwhelmingly by American corporations, as the tech chapter of this report shows. The Chinese state, for its part, directs capital through its state funds on a scale that this private statistic does not capture at all, an estimated 184 billion dollars cumulatively since the turn of the millennium.⁹⁹ What this gap feels like inside a single company is demonstrated by the SAP case described in the first chapter, because Europe's largest software group is financing its AI offensive not out of the strength of a capital market but out of the hiring, travel and supplier budgets of ongoing operations. At first glance these figures are sobering for Europe, yet Nagel himself expressly does not conclude from them that Europe must capitulate, and instead arrives at the assessment that the industrial AI race is open, because European and especially German industry can build on a wealth of data for training highly specialized models, with vendors such as Mistral AI (mistral.ai), Black Forest Labs (bfl.ai) or Aleph Alpha (aleph-alpha.com) and a whole series of barely known but internationally leading hidden champions in the background.⁶⁶

“SAP, Europe's largest software group, is financing its AI offensive not out of the strength of the capital market but through savings in ongoing operations.”

The difference from the usual European discussion is telling. Where sovereignty is mostly negotiated in public as a question of data protection and regulation, Nagel argues economically, namely that Europe's real raw material is not sheer computing power, in which it is structurally behind, but the depth and quality of industrial data, in which the continent has a genuine lead precisely because of its old, broad industrial base, if it knows how to use it. Even the productivity estimate Nagel cites, according to which a rapid diffusion of AI could bring the G7 states between eight tenths and one point three percentage points of additional annual productivity growth over the coming decade, comes from the most favorable of three scenarios in the underlying OECD study, whose full range across all three scenarios runs from 0.4 to 1.3 percentage points, a nuance that Nagel selectively omits. The condition under which the estimate holds at all matches the diagnosis

of this report, namely that organizations actually have to learn to work with the technology instead of merely buying it.⁶⁶

How hard European policy finds it to raise this raw material where it actually lies is shown, of all things, by the answer Brussels has found to its own competitive weakness. The Union's financial framework for the years 2028 to 2034, proposed in July 2025 and at just under 2 trillion euros the largest in its history, shifts the weights clearly toward competitiveness, defense and strategic autonomy; the new competitiveness fund alone, together with the doubled research program Horizon Europe, bundles around 409 billion euros, of which a good 51 billion is for digital leadership. The already mentioned paper "A Union under siege, or a Union besieging itself?" by the regional economists Andrés Rodríguez-Pose and Federico Bartalucci, who write as committed advocates of European cohesion policy and whose thrust must be read accordingly, nevertheless demonstrates that this architecture, conceived along sectoral and national lines, will with high probability direct its capital to where innovation systems are already dense, to Île-de-France, Upper Bavaria, Stockholm or Hamburg, and that the same few metropolitan regions will at the same time gather the defense and the autonomy gains, while cohesion policy, the only instrument that has so far reached the remaining regions, is dissolved into national plans. Should the proposal be adopted in this form, the European budget of the coming years will therefore reinforce precisely that split between capital corridors and the provinces which regional research has described for two decades as the real European pattern.¹⁰⁰

For the question of this report there is a double point in that. The industrial depth of data on which Nagel's confidence rests sits precisely not in the metropolitan corridors favored by the proposed financial framework, but in the mid-sized industrial regions in between, where the hidden champions have their plants and their process data and where at the same time hardly any of the new funding windows reaches. But if the European capital of the coming years flows above all to where strong ecosystems already stand, then the combination of artificial intelligence and a company's own domain expertise will not arrive in these regions as a program, neither from Brussels nor from the respective capital. Building this capability therefore remains, once again more sharply, a task for the individual company, and the sovereignty the European discussion talks about is decided less in the funding architectures than in the question of whether the companies of the industrial middle can turn their lead in data into productive practice themselves.

Alongside capital and education, a third axis of the sovereignty question is currently opening up, governance, that is, the question of who writes the rules by which the coming models are measured. How far from Europe these rules are taking shape is shown by a proposal that Demis Hassabis, Nobel laureate and the most prominent European born mind in top level AI research, made public in July 2026, at the time still as head of Google DeepMind, since August 2026 as chief scientist of the parent company Alphabet and chairman of the board of Google DeepMind, to be read as the positioning of a leading market player on its own behalf, not as neutral analysis. Modeled on the American financial regulator FINRA, he sketches a largely industry financed testing body that would examine the most capable AI models for dangerous capabilities before release, whose review would later become a condition for access to the American market, expressly independent of the model's country of origin, and which in an emergency could initiate a coordinated slowdown of development. Should this path prevail, the binding standards of AI development would arise in an American led, industry borne institution, and on this third question of

its sovereignty Europe would again stand only at the margin, as it already does on capital and on education.¹²⁴

On the question of how sound such self regulation by the industry would be, there has been an independent answer since the beginning of July 2026, though from an organization that is itself a party to the matter. The Future of Life Institute, an American organization founded in 2014 that advocates strict regulation and a slowdown of AI development, organized the 2023 call for a development pause and owes its assets largely to a donation from Ethereum founder Vitalik Buterin, had a seven member expert panel grade nine leading model vendors for the summer edition of its AI Safety Index, and the verdict is sobering. Not a single vendor gets beyond a C+, which in the American grading system corresponds to a bare satisfactory, roughly a three on the scheme customary in Germany; the leading vendors have weakened or in effect abandoned their earlier commitments to pause on reaching red lines; and at Google DeepMind of all places the panel points to the gap between the reassuring public communication of the leadership and the political conduct of the group. The panel's conclusion is that the credibility of safety precautions for the most capable models can no longer rest on self regulation but needs external verification, and this in the same month in which Hassabis proposes an industry financed form of self regulation. For Europe the index holds a point of its own, which the panel itself calls the European dissonance. While the Union leads the world in AI safety regulation, Mistral, the most important European AI company, lands in last place among all nine vendors examined, though Mistral has publicly disputed the evaluation logic. The grades are averaged expert judgments on the basis of public information, not audits, yet the finding fits the diagnosis of this chapter, because here too the ability to build safe models proves to be an achievement of the organization, not a result of the regulation that surrounds it.¹²⁵

How deep the European gap actually runs is shown by a look at the universities, and it also shifts the familiar picture of the triangle of AI powers. In university research China is by now at least on a par with the United States; according to the Stanford AI Index it leads in AI publications with a good 23 percent of the world share, in citations with just under 23 percent and in granted patents with almost 70 percent, and according to a Fortune analysis from November 2025 Tsinghua University in Beijing has produced more of the 100 most cited AI papers than any other university in the world. The American lead in frontier models, by contrast, rests not on its own young talent but on immigration. According to MacroPolo's Global AI Talent Tracker, 59 percent of the world's top researchers work at American institutions, yet most of them were trained elsewhere, above all in China, where 38 percent of the top researchers earned their first degree and only 11 percent have stayed in the country. Europe's position in this triangle is contradictory. Its universities are not the problem; their share of the world's cited AI papers, at just under 21 percent, is even ahead of the American 13 percent, yet the continent subsequently loses those it has trained to American industry, which according to industry reports pays 30 to 70 percent higher salaries for comparable roles, and it hardly converts its research into models and firms of its own.¹²⁶ The European deficit therefore does not begin in the lecture hall but behind it, in holding on to people and in making use of the knowledge, and it is thus the same capability question this report diagnoses at the level of companies, only one level higher. How far Europe has already been pushed to the edge in the American imagination is shown, incidentally, by the leading future scenarios of the US discussion themselves, because both the scenario AI 2027, read millions of times, in which a group around

former OpenAI researcher Daniel Kokotajlo played out the development month by month up to a superhuman AI in April 2025, and the policy draft AI 2040 added by the same group of authors in July 2026 conceive the coming order of artificial intelligence as a purely bilateral arrangement between Washington and Beijing, in which Europe no longer appears even as an actor.^{127, 128}

“59 percent of the world's top researchers work at American institutions, but most of them were trained elsewhere, above all in China with 38 percent.”

“Europe's universities provide just under 21 percent of the world's cited AI research, more than the American ones, yet the continent subsequently loses those it has trained to American industry.”

“The leading American future scenarios conceive the coming order of artificial intelligence as purely bilateral between Washington and Beijing; Europe does not appear in them even as an actor.”

“The European deficit does not begin in the lecture hall but behind it, in holding on to people and in making use of the knowledge.”

This closes the circle back to the demographic argument of the previous section. Europe has fewer workers but more industrial data than its competitors would have one believe, and which of the two forces prevails in the end will be decided not by the number of models but by the speed at which European companies actually turn the knowledge they have into productive practice, the question this report as a whole comes down to.

Alongside capital, education and governance, the sovereignty question is decided on a fourth axis as well, the political economy of acceptance, that is, the question of who in an economy has a tangible interest in the technology succeeding and who in its failing. The sharpest formulation of this question so far comes from Julia Willemyns, cofounder of the British think tank Centre for British Progress and thus herself a voice of the technology friendly camp, in an essay from July 2026 on the new Luddites. Her observation starts where the comparison with the Luddites of the early nineteenth century usually ends. The English weavers who destroyed the new looms in 1812 were put down by the military, because they had no political voice and the English elites profited from industrialization as a bloc. This time, by contrast, the upheaval hits the occupations of the well organized middle, from law firms through newsrooms into public administration, that is, groups that have professional chambers, associations and voting power behind them, and by Mancur Olson's classic logic the well organized groups on which the costs are concentrated regularly prevail in a democracy against the many across whom the benefits are spread thin. For the United States, Willemyns nevertheless expects that the resistance can grow loud and still lose, because capital markets, the federal budget and geopolitics there are by now so tightly coupled to the success of the technology, down to direct federal stakes in key firms, that the country's governing interests hang on its succeeding. Europe lacks exactly this counterpart: there is no European equivalent to

the large technology groups and no treasury that would share appreciably in the AI surplus and could compensate the losers of the rebuild out of it, but there is a decision making architecture full of veto points. On this reading, the capital gap with which this chapter began would be not only an innovation problem but also an acceptance problem, because where no gain of one's own arises there is nothing to distribute, and the political space belongs to those who have something to lose.¹²⁹

“The weavers were put down because the English elites profited from industrialization as a bloc. This time the upheaval hits occupations that are well organized and have associations and chambers behind them.”

“Europe lacks the counterpart: no large technology group, no treasury that would share appreciably in the AI surplus in order to compensate the losers of the rebuild.”

What does not appear in this calculation, however, is demography, and with it the account lacks precisely the factor that actually distinguishes Europe's political economy from the American one. A labor market that, by the figures in this part, is running out of people simply produces fewer losers to be compensated than the blockade thesis assumes, and an employment system that legally pushes toward retraining rather than separation turns potential objector groups into participants in the rebuild. Whether Europe has the political economy most prone to blockade or the one most capable of consensus for this technology is therefore not a settled matter, but depends on whether the introduction of digital team members in the company is experienced as an answer to the shortage of workers or whether the shortage merely serves as a pretext for it. In the end this too is not a question of regulation but of leadership in the companies themselves.

Looking ahead

If one puts together the building blocks of this part, the missing collapse in the Union's aggregate figures, the separate controversy in the individual countries, dismissal protection and demography, then a possibility emerges for Europe that does not appear at all in the imported American discussion. Europe could become the first continent to introduce artificial intelligence not as an instrument of downsizing but as an answer to a labor market that is running out of people. The ingredients for that are in place: an employment system that legally pushes toward retraining rather than separation, an industrial base of data whose depth the competition cannot copy, and a demographic pressure under which any productivity raising technology no longer appears as a threat scenario but as a promise of relief.

“Europe could become the first continent to introduce artificial intelligence not as an instrument of downsizing but as an answer to a labor market that is running out of people.”

What is missing is not further regulation and not a European frontier model either, but the organizational ability to bring these ingredients into ongoing operations, that is, the same ability which, by everything this report shows, even the most capital rich are only now working on. In this one respect the race does not begin for Europe from behind, but with an open outcome.

Nine predictions for 2027 that this report will grade publicly next year

This chapter is the part of the report that turns it from a snapshot into a series. It contains nine predictions about the coming year, each worded so that in the fall of 2027 it can be said without ambiguity whether it came true, each with the measure against which it will be tested, and each as a personal bet by the author, not as the output of a model. The 2027 edition grades them publicly, counts the hits and the misses, and explains what was wrongly reasoned in the misses. Alongside them, the chapter keeps a block of external benchmarks, meaning already dated predictions by vendors, market researchers and researchers on the same subject, with the source, the target year and the stated self-interest. They do not replace a prediction of one's own, but they will be graded along with the rest in the following year, because nobody learns anything from a debate in which new forecasts are made every year and the old ones are forgotten. Where an external prediction names no measurement criterion, that is noted, and the grading is tied to observable indicators that the report itself defines.

The predictions follow the order of the report, from the loud first force to the organization. They are the personal view of the author, formed from the findings of this report, and they state what he himself, given the evidence, considers the most likely course of the coming year. They are deliberately cautious where the evidence is thin, and deliberately sharp where the report commits itself. Anyone who thinks they are wrong will find in each one the place where that will show in 2027.

First prediction: the AI argument remains a front rank reason for layoffs in the United States in 2027, while the officially measured effect stays small. In the layoff statistics of Challenger, Gray & Christmas, the technology will again be among the three most frequently cited reasons in at least six months of 2027, and at the same time fewer than 10 percent of AI using companies in the business survey of the American Census Bureau will continue to report a reduced headcount. The prediction is a bet that the story and the measurement keep drifting apart, because the AI argument is an instrument of capital market communication and not a description of the operating business. It will be tested against the monthly Challenger reports and against the next AI supplement to the Business Trends and Outlook Survey.

Second prediction: at least three well known companies publicly reverse an AI justified hiring freeze at the entry level by the fall of 2027. The thinned out pipeline of their own junior talent and the rising cost of experienced specialists force the change of course, and it will not be carried out quietly but announced as a program, as Cognizant and EY already demonstrated in 2026. It counts as a hit when a company with more than 10,000 employees publicly announces that it is explicitly withdrawing or reversing an earlier reduction of entry level positions that was justified by the technology; programs that were never preceded by a freeze do not count.

Third prediction: the American entry level slump will again not appear in the official series of Germany and Austria in 2027, and the delay persists. Unemployment among people under 25 with a

university degree will not rise more steeply than overall unemployment in either country, and the continuation of the Swiss KOF measurement remains the only European finding in which exposed occupations perform measurably worse. The prediction is the delay hypothesis from the chapter on the European yardstick in testable form; it is refuted if the Bundesagentur für Arbeit or the Arbeitsmarktservice reports a disproportionate increase among young graduates in exposed occupations for 2027.

Fourth prediction: the gap between pilot and production does not close in 2027. In McKinsey's annual survey series, the share of companies that attribute an effect on the operating result to the technology will stay below 45 percent, after 37 percent in 2026, a value that McKinsey itself reports as unchanged from the prior year. The group with a measurable contribution to operating profit of at least 5 percent will stay below 10 percent, after around 6 percent. Adoption will keep growing, the effect will not grow to the same degree, because organizations learn more slowly than the vendors deliver. It will be tested against the 2027 edition of the same series, with the caveat that these are self reported answers from a survey run by a consultancy that sells rebuild programs.

“Organizations learn more slowly than the vendors deliver, and the group with a noteworthy contribution to operating profit stays below 10 percent in 2027 as well.”

Fifth prediction: at least five well known companies in the German speaking countries announce a reorganization by the fall of 2027 that structures units around topics and outcomes instead of functions. The organization of the future, described this year as a target picture by consultancies, market researchers and tool vendors, reaches practice first in companies with a high share of knowledge work and light regulation, and it will be announced publicly there, because it can be marketed as a competitive advantage. Announced reorganizations count as a hit when small units take over responsibility for the entire process, as do the corresponding answers in the 2027 board pulse, into which the question about the division of work will be taken up as a core question.

Sixth prediction: at least one publicly visible large scale attempt to flatten hierarchies and align the organization around AI agents will be partially withdrawn by the fall of 2027. Meta's Project OT is the first documented case in 2026, and it will not be the last, because in most cases the coordination work that was cut turns out to have been quietly shifted rather than saved, following the pattern of Zappos and Medium. It counts as a hit when a company with more than 10,000 employees publicly withdraws or suspends an announced reduction of hierarchy; a mere slowdown does not count.

Seventh prediction: the gap between productivity and wages widens further in the early years of the technology. In the American series of the Economic Policy Institute and the Bureau of Labor Statistics, productivity will grow more strongly in 2026 than the hourly wage of typical employees. In the German speaking countries, the spread between entrants and experienced staff in companies' wage expectations will not narrow if the survey is repeated; the ifo Institute measured this spread for the first time in 2026. The prediction puts the distribution question of the closing chapter into testable form, and it is refuted if wages keep pace with productivity in 2026.

Eighth prediction: the AI argument does not arrive in the German speaking countries in 2027. No company in Germany, Austria or Switzerland with more than 10,000 employees justifies job cuts of more than 1,000 positions primarily with artificial intelligence by the end of 2027. What has been part of capital market communication in the United States since 2025 fails to appear in the German speaking countries, and not because the technology is used less here, but because the argument yields nothing here. Where job cuts have to be negotiated with a works council, quantified in a social plan and justified before labor courts, invoking the technology is a poor justification, because it invites the question of which task exactly the machine has taken over. The job cuts that come in 2027 will be justified with the business cycle, costs, location and restructuring, the way SAP ran its restructuring of around 10,000 positions as a realignment and, in July 2026, saved on hiring, travel and suppliers in order to finance its own AI push, without attributing a single job to the technology, and the way Ingka retrained 8,500 call center employees instead of laying them off. It will be tested against the mandatory disclosures of listed companies, against press releases on restructurings, and against the social plans and mass layoff notifications of 2027 that become public. A counterexample counts only if the company itself names the technology as the main reason in its own announcement, not if the press attributes it to the company; a single counterexample refutes the prediction.

Ninth prediction: the gender gap in the professional use of generative AI does not close in 2026. In Eurostat's statistics on ICT use in households, men in Germany use generative AI at work 29 percent more often than women, 18 against 14 percent. In the annual figure for 2026, which Eurostat publishes in December 2026, this lead remains above 20 percent, even though women aged 16 to 24 were already ahead of men of the same age in 2025. The gap arises in the workplace, and in 2026 companies are not reorganizing their training fast enough for a single age cohort to close it. It is measured by the ratio of the professional usage rates of men and women between 16 and 74; the prediction is refuted if the men's lead falls to 20 percent or below. Austria runs alongside as a control series; there, the lead stood at 30 percent in 2025.⁶

The author considered two further predictions and did not include them, because their measure will not yet be cleanly available in 2027. According to the first, the midterm elections in the United States in November 2026 become the first election in which concern about AI driven job loss is a load bearing campaign issue. According to the other, at least two further large commerce platforms follow Shopify's example and report AI mediated orders as a separate metric. They stand here so that the 2027 edition can also check whether leaving them out was right.

The external benchmarks that this report archives and grades along with its own

SOURCE	PREDICTION	TARGET YEAR	SELF-INTEREST	STATUS IN SEPTEMBER 2026
Geoffrey Hinton, 2016 ⁷⁷	Training radiologists should be stopped, because machine learning would read images better than humans within five years.	2021	Researcher, no commercial interest	Missed. Ten years later, the official French statistics count more radiologists, more assistants and more imaging centers than in 2014, with bottlenecks in recruitment.

SOURCE	PREDICTION	TARGET YEAR	SELF-INTEREST	STATUS IN SEPTEMBER 2026
Dario Amodei, Anthropic, May 2025 ²⁷	Up to half of all entry level positions could disappear within five years.	2030	Vendor whose business depends on the effectiveness of the technology	Qualified by the source himself; reinterpreted in 2026 as a management decision. No measurement criterion named; the report tests against the hiring rate of 22 to 25 year olds in exposed occupations. First interim status: the unemployment rate of more than 10 percent announced alongside it has not materialized, it stood at 4.1 percent in August 2026, while hiring of 22 to 25 year olds in the most replaceable occupations has fallen by roughly a third since 2021. ^{136, 138}
AI Futures Project, AI 2027, April 2025 ¹²⁷	A system that completes every programming task faster and more cheaply than the best human developers, in the course of 2027; a near complete hiring freeze for programmers from mid 2027.	2027	Authors from safety research, not a vendor	Pushed back by the authors themselves; the median now lies around 2030. The scenario's labor market claims have no evidence base of their own.
Gartner, July 2026 ⁵³	60 percent of organizations move their software development to small teams across the board, after 15 percent in 2026.	2029	Market researcher, sells consulting on the rebuild it describes	Open. The methodology behind the 15 and 60 percent is not disclosed, an analyst estimate.
Gartner, July 2026 ⁵³	Companies that forgo entry level hires today will hollow out their own junior talent base by 2028.	2028	as above	Open. No measurement criterion named; the report tests against the share of entrants among new hires in exposed occupations and against the salary premiums for experienced specialists.
Zach Lloyd, Warp, July 2026 ⁹³	Within a year every significant software project runs an automated development engine; the share of automatically submitted code changes rises from around 20 to between 30 and 60 percent.	2027	Vendor of the tool described	Open. It will be tested against vendor statistics on automatically submitted code changes, with the caveat that these measure output and not value.
Goldman Sachs, 2026 ²⁹	With full diffusion of today's use cases, 2.5 percent of American employees are affected, the unemployment rate sits half a percentage point above trend during the transition phase and returns to it within two years.	2028	Investment bank with exposure to the technology industry	Open. So far the measured effect lies below this; the benchmark serves as an upper bound for the first force.
Anthropic Institute, September 2026 ¹³⁰	In the moderate and the substantial scenario, the American unemployment rate stands at 3.9 and 4.6 percent respectively at the beginning of 2030; only in the extreme scenario does it reach 17.9 percent in office and knowledge occupations.	2030	Research unit of a vendor	Not a prediction but a model calculation; the report carries it as a frame against which the job killer story can be measured, and in 2027 it tests the underlying assumptions about adoption.

The table grows with every edition, and no row is deleted. Anyone who puts a number into the world in 2026 will find it here again in 2027, next to what became of it.

The future of work is decided not by the machine but by the company's ability to work with it

This report began with the promise of breaking a headline down into its components, and at the end of that path the great misconception can be pinpointed exactly. The job killer story makes the machine the actor, and that is precisely what does not hold up against any of the three forces examined here. Behind the AI argument stand capital market logic and restructuring, for which companies borrow a technological justification. The career gap is real and measurable, yet even where the technology demonstrably plays a part, it hits a stage whose value companies determine through their own hiring practice. And the long rebuild, the actual effect of the technology, is stalling precisely not on the capabilities of the models but on the ability of organizations to put them to work in ongoing operations. For all three forces the investigation arrives at the same result: employment is decided by people and organizations, not by the autonomy of a technology.

How far from self-evident it is that the gains of this rebuild reach the people who do the work is shown by the longest series that labor statistics offer. Since the 1970s, productivity in the United States has grown considerably faster than real hourly wages, and the same pattern runs through the United Kingdom, Japan and large parts of Europe; the Economic Policy Institute puts the gap for the period since 1979 at a productivity gain of a good 92 percent against a gain of just under 34 percent in the hourly wage of typical workers. Should artificial intelligence further strengthen capital's share of value creation, there is much to suggest that this gap will widen rather than close; the Financial Times recently distilled this concern into the formula that the productivity gains of the technology could flow above all to the owners of the machines and data centers. Erik Brynjolfsson, whose Stanford lab measured the career gap, answers in the same piece with a distinction that confirms this study's central thesis from the economic side: whether the technology replaces people or complements them is a question of design and not a property of the machine. The trap he warns about he already called the Turing trap in 2022, the building of systems that imitate people instead of complementing them. For this study's question that means: the distribution of the gains, too, is not a natural constant of the technology but a consequence of decisions, in companies as in politics, and the rebuild will be measured by whether it not only raises productivity but also shows up in the wages of those who carry it.¹³¹ That this concern is no preserve of voices critical of capitalism is shown by the fact that the economics department of a German state bank explicitly assumed in August 2026, in all of its own scenarios, that without state intervention the lion's share of the higher labor productivity does not flow to employees in the form of rising wages but increases the profits of capital owners.³³

That this distribution question is not a distant one is suggested by an analysis that came out of the financial industry in July 2026 and has so far remained without a European counterpart. Torsten Slok, chief economist of Apollo Global Management (apollo.com), and the Princeton researcher Sania Edlich compared wage developments from 2015 to 2026 for 321 matched American occupations, and they determined exposure not theoretically but by the observed occupational use of Anthropic's models. In the highly exposed occupations, real wage growth since 2023 has been

6.7 percent lower than in the weakly exposed ones, while employment in the same occupations has remained unchanged, and the gap widens toward the bottom: in the lowest wage quartile it is around 11 percent, in the two middle quartiles 5.4 and 4.0 percent, and among top earners it is no longer detectable. Should this finding hold, the effect of the technology on the existing workforce shows up not in the job but in the price of labor, and it shows up first among those who earn the least. That is the same movement this chapter describes along the long series, only at the current edge and with the technology as the named cause, and it also supplies the series against which this report's prediction on wage developments can be checked next year. An honest reading includes the fact that the calculation here comes from the research department of an asset manager and that no peer reviewed work exists, that exposure is drawn from the usage data of a single vendor and the finding therefore must not be listed as an independent piece of evidence alongside that vendor's other analyses, and that the authors themselves leave open whether their values denote percent or percentage points, which is why the report adopts only the direction and the ranking by wage quartile.¹³⁷ It is notable that Benjamin Verschuere, one of the authors of the task decomposition cited earlier, arrives at this very question of his own accord in the same front page story when he notes that wage growth has become detached from productivity.¹³⁸ How far, the official statistics say: the labor share in the American nonfarm business sector, that is the share of value added that goes to employees as compensation, stood at 52.8 percent in the second quarter of 2026, lower than at any time since this series began in 1947, after a peak of 66.2 percent at the end of 1960.¹³⁹ That is no work of artificial intelligence, because the decline has been running since the early 2000s, but it is the starting position into which this technology acts, and it explains why the question of how the gains are distributed cannot be postponed to the day the rebuild is complete.

In September 2026 an AI vendor itself ran the numbers on how large the upheavals in the labor market could turn out to be. In a working paper from the Anthropic Institute, the company's research arm, the economists Anton Korinek and Charles Jones, together with further authors, sketch three scenarios for the American economy through 2030, explicitly without stating probabilities. The finding is simple. In the two scenarios with moderate and with substantial change, the unemployment rate in 2030 lies between 3.9 and 4.6 percent, that is, close to today's level, and the model arrives at the same values if, instead of the authors' assumptions, one enters the expectations from a representative survey of American adults. The image of the job killer emerges only in the extreme scenario, in which the economy grows by 15 percent a year, artificial intelligence handles nine out of ten work steps on its own and no new tasks arise for people; unemployment in office and knowledge occupations then rises to 17.9 percent. What does shift as early as the middle scenario is the distribution: the share of labor income in all income falls from 60 to 56 percent by 2030, and in the extreme case the entire increase flows to capital income. The model applies only to the United States, leaves political countermeasures out and rests in part on the vendor's own usage data, by the standard already presented in the chapter on the career gap.⁴⁰ The image of the job killer requires the most extreme assumptions of this calculation; for an unequal distribution of the gains the middle ones already suffice.¹³⁰

How far off the job killer story is shows most clearly where its proponents do not look at all, at new company formations. The study by Kim and Koning, which has already supplied evidence at two points in this study, contains a third finding as well. AI-native new companies are around a quarter smaller and about half a hierarchy level flatter than comparable firms in the same industry and the

same founding cohort, and in doing so they reach comparable or higher valuations. One can read this as confirming the fear of shrinking workforces, but that misjudges what is actually happening at these companies. No work is disappearing here; new value creation is arising in a different organizational form, smaller in setup, staffed differently and geared from the start to the capabilities of digital team members. The work does not disappear, it changes place and shape, it migrates into products, into new firms and into roles that do not appear in the personnel statistics of existing companies. Anyone who only counts which positions are cut at established companies is taking stock of an economy's rebuild with the categories of its old state.⁵¹

That leaves the question this report put at the start and carried through every chapter: not whether artificial intelligence destroys jobs, but how fast an organization builds the knowledge with which it makes the technology productive. The evidence that the fault line runs exactly here has grown denser from chapter to chapter. The gap between pilot and production has been measured by three mutually independent findings. The delegation experiment by Colaiacovo and Koning demonstrates that the willingness to hand work over to an AI actor depends only to a good third on what one credits the technology with, and beyond that is a management decision in its own right.⁷ The models themselves can by now be rented by any company, at falling prices and in increasingly interchangeable quality. The ability to work with them cannot be rented.

“Any company can rent models. The ability to work with them every company has to build for itself.”

This diagnosis no longer stands alone. In June 2026 the Stanford Institute for Human-Centered Artificial Intelligence presented an industry report that was produced within the institute's corporate-funded membership program and that condenses a research overview by more than 50 authors; its central message coincides almost entirely with this study's central thesis. Because the large foundation models are converging in performance, the difference between vendors shifts from the question of what the model can do to the question of how well it serves the people who work with it. Top scores in the usual performance tests accordingly guarantee no business benefit, productivity gains arise only when employees actually change the way they work, and the design of the collaboration between a human and an AI colleague often brings greater progress than switching to the next model. That of all places the academic home of human-centered AI research locates the competitive advantage where this report diagnoses the learning gap shows that the central thesis is no European peculiarity but reflects the state of the international discussion.⁹⁵

International management scholarship has by now arrived at this point as well. Gary Hamel and Michele Zanini, whose work on organizational research has shaped the discussion for decades, measured the industry's three loudest promises against the evidence in June 2026, mass displacement, imminent general machine intelligence and the productivity revolution, and arrive at the same finding as this study: the evidence does not support the headlines, the gains are real where tasks are clearly delineated. Their most robust figure comes from the same calculation by the Penn Wharton Budget Model with which it was already shown above what condition the European hope for higher productivity is tied to. What is remarkable is less the agreement than the open flank, because they conclude from it that expectations must be recalibrated, and they attribute part of the gap, unlike this study, to the limits of today's models themselves. This study's

findings go one step further: the gap between promise and effect is no fate of the technology but the distance between its capabilities and the ability of organizations to use them. That the most prominent critics of the hype arrive at the diagnosis and stop short of this explanation underlines why this report puts the capability question at the center.²⁹

A sector-specific confirmation comes from the very industry that traditionally proceeds most cautiously with new technology. The study by Wavestone and the University of St. Gallen already cited in the chapter on the long rebuild, a survey of 30 business unit and IT managers from banks, insurers and technology vendors in the German-speaking region, draws a conclusion that points beyond the financial sector: what decides how far a company can actually scale agentic AI is not the capability of the models but the maturity of its operating model. That such a regulated and risk-aware industry arrives at the same conclusion as management research and academic AI research before it shows that this report's central thesis describes no isolated case of the long rebuild but a pattern that repeats across industries and continents.⁶⁸

How far this fault line reaches beyond the individual company is shown by an essay that the economist E. Glen Weyl, the sociologist James Evans and the Microsoft research director Chris White published in July 2026 in the magazine *Noema* together with a 23-member coalition from research, business and civil society; two of the three authors do research for Microsoft, and the text nevertheless criticizes the investment logic of their own industry. For every dollar that flows into more capable models and their alignment with people, only a fraction goes into the ability of institutions, organizations and education systems to absorb these systems, and the authors call this neglected counter-direction reverse alignment. If it fails to come, three wrong turns loom, which the authors call productivity without prosperity, execution without review and capacity without limits, and the first of them is the same distributional flank that became visible above in the wage-productivity gap. Their diagnosis for the world of work repeats almost word for word the lesson this study drew from the introduction of the electric motor, since the real bottleneck for large productivity gains is not the model but the rebuild with which organizations align themselves to the possibilities of the technology. What this report describes as the capability of the individual company returns there as a task for entire societies, and both levels condition each other, because no company builds its capability more easily than in a country whose education, labor market and legal institutions support the rebuild.¹³²

What such a national counter-direction could look like has been on view in the United Kingdom since the start of the year. There, the University of York, the higher education data provider QS and the consultancy Public First presented in January 2026 the attempt to quantify the education side of the capability question for an entire national economy, along the eight priority sectors of the British industrial strategy. Of the 1,890 occupations into which the underlying market model divides the British labor market, the analysis classifies 1,436 as critical for exactly one of these sectors and a further 635 as critical for several at once, whose scarcity would therefore slow several growth strategies at the same time, and more than 80 percent of the most important of these occupations require a university degree. The same model estimates the complementing potential of the technology to be considerably larger than its displacement potential, since three quarters of employees work in occupations whose work the technology can at least noticeably complement, while only half are exposed to automation to a comparable degree. The report's conclusion transfers this study's central thesis to the entire national economy, since employment growth will in

future be less a question of demand in individual industries than the question of whether the workforce masters those tasks that the technology complements instead of replacing, which makes higher education and lifelong learning systemically critical infrastructure. An honest reading includes the fact that parties with interests of their own are doing the calculating here, a university in the funding crisis of its sector and a commercial higher education data provider, and that the most spectacular figure of the report, an AI growth contribution of 490 billion pounds by 2030, rests on the heroic assumption of up to two AI-supported working days per week and thus lies many times above the cautious productivity estimates with which it was shown above what condition the European hope is tied to. The basic statement holds even without this set of figures, however, since the government agency Skills England expects independently of it that two thirds of the additional employment in the priority occupations will require a qualification above school education, and precisely there lies the gain for this study's capability question: a country has begun to look for the bottleneck of its growth no longer in the technology but in the education of the people who are to work with it.¹³³

By now it can also be read off how employees themselves experience this fault line. The second annual sentiment survey by Lenny's Newsletter among 5,332 working people in the technology industry, to be read as a self-selected mood picture strongly shaped by product management and by American respondents and not as a labor market measurement, describes a workforce that splits into two camps. 49 percent experience themselves as strengthened by the technology in their professional role, just under 19 percent as unsettled or devalued, and this self-placement predicts career optimism and the willingness to recommend one's own field more strongly than role, hierarchy level and company size together; for optimism, the distance between the strengthened and the devalued corresponds to about three times the effect that a founder role brings with it. At the same time, the share of respondents with pronounced burnout rose within one year from just under 45 to a good 55 percent. So the same question that divides organizations into learners and onlookers divides working lives as well, and what tips the scale is the experience of whether the technology carries you or runs you over. An organization that builds its capability thereby also decides on which side of this line its people stand.²⁸

What this capability consists of at its core was already named in the chapter on the long rebuild; here the findings are brought together. Between the model and one's own process lies what matters: the company's knowledge, its rules, its experience and its decisions, held in a form that belongs to the company itself and that digital team members can work with directly, no matter which model happens to underlie it. Whoever owns this foundation can switch vendors like tools and takes their knowledge along every time. Whoever does not own it hands a piece of knowledge work to their supplier with every AI function they buy and calls that progress. The decision as to which work may migrate into the purchased product and which remains tied to one's own processes is made task area by task area and is therefore not a software question but the form in which a company's capability becomes concrete. It is remarkable that an entirely different investigation arrives at the same order of magnitude. When a research group from the Massachusetts Institute of Technology and Stanford University broke down the differences in the advice that language models give private individuals about their finances, there too only a third was attributable to the model itself and two thirds to how the people had asked.⁸¹

How seriously this shift is now being taken is shown by the supply side of the model market itself. Thinking Machines Lab (thinkingmachines.ai), the company of former OpenAI technology chief Mira Murati, presented a manifesto in July 2026, a few days before the release of its first model, that locates its own business foundation precisely here. Organizations should train the models further with their own knowledge instead of buying an artificial intelligence that skims this knowledge off once and replaces it with a standard offering. That is a vendor text in its own cause and is to be read as such, yet when a vendor of top models looks for the future of its business where companies keep their knowledge in their own operations, the market is betting at this point on the same capability this study is about.¹³⁴

For Europe this calculation turns out sharper, for better and for worse. The second part of this study has shown that the starting position is more favorable than the imported discussion suggests: a demography in which any technology that raises productivity does not appear as a threat but promises relief, an employment system that pushes toward adaptation instead of separation, and an industrial depth of data that sits in the mid-sized regions and cannot be copied. But all three advantages are conditional. They pay off only for organizations that build the capability to use them, and this work, as the look at the European financial framework has shown, is taken off their hands by no funding scheme. The location supplies the ingredients, but every company has to process them itself.

Looking ahead

When this assessment appears for the second time next year, it will ask the same question again, with the same core questions to the same group of decision makers, with the same metric for the distance between pilot and production and with the public grading of its own predictions. Only this recurrence turns a diagnosis into a series of measurements, and only the series of measurements will show whether the learning gap is closing or whether it is opening further month by month, like the American figures on the career gap.

The great misconception about the AI job killer will settle itself in the process, as technological figures of fear have always settled themselves, not through refutation but through precision. What remains is the question on which the coming years will be decided: which organizations have made productive digital colleagues out of a capable tool, which have rebuilt the entry level, because tomorrow's capability too needs entrants who are allowed to learn today, and which no longer need the AI argument because their numbers speak for themselves. To the question of whether artificial intelligence destroys work there is, at the end of this study, a precise answer: it depends on who leads it. That is not a reassurance but the most exact description of the situation that this year's evidence permits, and it places responsibility where it belongs, in the organizations themselves.

Sources

- 1 Carlota Perez: "Technological Revolutions and Financial Capital. The Dynamics of Bubbles and Golden Ages". Edward Elgar, Cheltenham 2002.
- 2 Kathryn Bonney, Cory Breau, Emin Dinlersoz, Lucia Foster, John Haltiwanger, Aditya Pande: "The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks". CES Working Paper 26-25, Center for Economic Studies, U.S. Census Bureau, April 2026, also NBER Working Paper 35141. <https://www.census.gov/library/working-papers/2026/adrm/CES-WP-26-25.html>, PDF at <https://www2.census.gov/library/working-papers/2026/adrm/ces/CES-WP-26-25.pdf>
- 3 Dan Tinkoff, Lieven Van der Veken and Michael Chui, with Tara Balakrishnan: "The state of AI in 2026: On the road to ROI". QuantumBlack, AI by McKinsey / McKinsey & Company, August 25, 2026. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>.
- 4 Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence. Erik Brynjolfsson, Bharat Chandar, Ruyu Chen. Working Paper, Stanford Digital Economy Lab / SIEPR, authoritative version of August 12, 2026 with ADP data through June 2026 (earlier versions of November 13, 2025 and of August 26, 2025). <https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence/>, PDF at https://digitaleconomy.stanford.edu/app/uploads/2026/08/Canaries_August2026.pdf
- 5 Apprenticeship market figures from the same Deutschlandfunk interview of August 15, 2026, see³⁴, reported by Evangelischer Pressedienst on August 15, 2026, <https://www.evangelisch.de/inhalte/258175/15-08-2026/nahles-162000-ausbildungsstellen-noch-unbesetzt>.
- 6 Eurostat, Statistical Office of the European Union: Survey on ICT usage in households and by individuals 2025, database table isoc_ai_iaiu (Use of generative AI tools by individuals), published December 16, 2025. https://ec.europa.eu/eurostat/databrowser/product/page/ISOC_AI_IAIU; the relative gaps are calculated from the rates in the table. As well as Initiative D21 e. V. and Institut für Arbeitsmarkt- und Berufsforschung (IAB) der Bundesagentur für Arbeit: "Digital Gender Gap. Schwerpunkt 2026: Künstliche Intelligenz", April 2026. https://initiated21.de/uploads/03_Studien-Publikationen/Digital-Gender-Gap/D21-GenderAIGap-2026.pdf, along with the press release of the IAB, <https://iab.de/presseinfo/gender-ai-gap/>, and the article in IAB-Forum, <https://iab-forum.de/der-gender-ai-gap-ki-wird-zur-schlueselressource-aber-maenner-und-frauen-nutzen-sie-nicht-gleich/>.
- 7 Colaiacovo, I. and Koning, R. (2026). Delegating to AI: Beliefs and the Organization of Work at New Ventures. Harvard Business School Working Paper, No. 26-093, June 2026 (PDF metadata date the full text to June 18, 2026). hbs.edu/faculty/Pages/download.aspx?name=26-093.pdf
- 8 Shaw, S. D. and Nave, G. Thinking, Fast, Slow, and Artificial: How AI is Reshaping Human Reasoning and the Rise of Cognitive Surrender. The Wharton School, University of Pennsylvania, SSRN Working Paper No. 6097646, posted February 2, 2026.
- 9 Dell'Acqua, Fabrizio, Edward McFowland III, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Kraymer, François Candelon and Karim R. Lakhani. Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality. Harvard Business School Technology and Operations Management Unit Working Paper No. 24-013, September 18, 2023 (revised September 27, 2023); published peer reviewed in Organization Science (2025), DOI: 10.1287/orsc.2025.21838.
- 10 Work at the Frontier: How AI is expanding what people do at work. Caroline Chin and Alex Martin Richmond, OpenAI Economic Research, July 2026 (published July 27, 2026). <https://openai.com/index/how-ai-is-expanding-what-people-do-at-work/>
- 11 Julian Kirchherr, Karel Eloit, Sandra Durth, Ulf Schrader and Vincent Bérubé, with Charlotte Seiler, Kira Rupietta, Kristina Störk, Marlene Senst, Séverine Fobe and Simon Gallot Lavallée: "HR Monitor 2026. A turning point for the people function". McKinsey & Company, People & Organizational Performance Practice, June 8, 2026, 44 pages. <https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/hr-monitor>.
- 12 Asmus Komm, Bryan Hancock and Sandra Durth, with Christian Winnewisser: "Escaping the pilot trap: Building HR for the agentic era". McKinsey & Company, People & Organizational Performance Practice, August 4, 2026, 13 pages. <https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/escaping-the-pilot-trap-building-hr-for-the-agentic-era>.

- 13 Sir Martin Sorrell (Executive Chairman, S4 Capital, previously founder and longtime CEO of WPP), LinkedIn post of July 24, 2026, announcing an episode of the podcast The Business of Marketing with Justin Cooke, recorded at the Cannes Lions. https://www.linkedin.com/posts/sir-martin-sorrell-35279a67_the-industry-is-focused-on-the-wrong-thing-share-7486440504973250560-NIKW/
- 14 Advertising market and platform figures: WARC, The Future of Commerce Media 2025 (November 4, 2025) as well as Global Ad Spend Forecast (Q2 2025 and December 2025): retail media 174.9 billion US dollars in 2025, 196.7 billion in 2026, share of the advertising market 14.7 to around 16 percent depending on the report; Alphabet, Amazon and Meta together around 55 percent of the advertising market outside China. GroupM/WPP Media, This Year Next Year (December 2024): 176.9 billion for retail media in 2025. Amazon, Earnings Release Q2 2026 (July 30, 2026): advertising revenue 19.8 billion, up 26 percent. Meta, Earnings Call Q2 2026 (July 30, 2026): advertising revenue 59.4 billion, up 27 percent. Zuckerberg statement: Stratechery interview with Ben Thompson, May 1, 2025. Google Ads Blog, April 15, 2026 (AI Max migration timeline), confirmed by Search Engine Land, August 14, 2026. Marketing Brew, April 7, 2026.
- 15 Mandatory disclosures and earnings communications of the agency holding companies: WPP, Preliminary Results 2025 (February 26, 2026) and Interim Results 2026 (August 6, 2026), profit warning of July 9, 2025 with a share price loss of 16 percent (Bloomberg, Forbes); Publicis Groupe, full year results 2025 (February 3, 2026) and half year results 2026 (July 16, 2026); Omnicom, completion of the merger with Interpublic (November 26, 2025) and Q2 results 2026 (July 28, 2026); S4 Capital, Preliminary Results 2025 (March 23, 2026) and Interim Results 2026 (August 5, 2026).
- 16 ANA, In-House Agencies No Longer A Trend (survey February/March 2023, 162 respondents, five year series since 2008, based on ANA members); Gartner, CMO Spend Survey 2025 (press release May 12, 2025, 402 CMOs); The CMO Survey of Duke University, 35th wave (Christine Moorman, survey January 2026, 308 respondents); Klarna, press release of May 28, 2024; Unilever, company announcement on the Sketch Pro studios, July 2025; ISBA/CvE, State of In-Housing 2023, single sourced only and carried with that qualification.
- 17 Statistisches Bundesamt (Destatis): "13,3 Millionen Erwerbspersonen erreichen in den nächsten 15 Jahren das gesetzliche Rentenalter", press release PD26_N039, June 2026, according to which by 2040 around 13.3 million people in the labor force, and thus 30.0 percent of all people in the labor force, will pass retirement age. https://www.destatis.de/DE/Presse/Pressemitteilungen/2026/06/PD26_N039_13.html.
- 18 Gartner survey of 350 large companies, each with more than one billion US dollars in revenue, on job cuts and the economic success of AI investments, spring 2026, with commentary by Helen Poitevin (Gartner Research).
- 19 Marc Andreessen (Andreessen Horowitz) on the podcast 20VC by Harry Stebbings, episode of March 30, 2026, on overstaffing at large companies: "Essentially, every large company is overstaffed. It's at least overstaffed by 25%. I think most large companies are overstaffed by 50%. I think a lot of them are overstaffed by 75%", and on the AI argument: "now they all have the silver bullet excuse, right? Ah, it's AI." <https://podcasts.apple.com/us/podcast/20vc-marc-andreessen-on-the-future-of-venture/id958230465?i=1000758169390>, confirmed secondarily by Fortune, March 31, 2026. Sam Altman (OpenAI) in the CNBC interview on the sidelines of the India AI Impact Summit, New Delhi, February 19, 2026, exclusively on the AI argument: "there's some AI washing where people are blaming AI for layoffs that they would otherwise do, and then there's some real displacement by AI of different kinds of jobs." <https://www.cnbc.com/video/2026/02/19/openai-sam-altman-india-ai-summit-full-interview.html>
- 20 Jessica Liu and Douglas Webber. "AI Adoption and Firms' Job-Posting Behavior". FEDS Notes, Board of Governors of the Federal Reserve System, Washington D.C., March 27, 2026. DOI 10.17016/2380-7172.4026. <https://www.federalreserve.gov/econres/notes/feds-notes/ai-adoption-and-firms-job-posting-behavior-20260327.html>.
- 21 Ara Kharazian (Ramp), Lisa Simon (Revelio Labs) and Ryan Stevens (Ramp): "A New Look at AI's Impact on Jobs: Firm-Level AI Spending and Workforce Adjustment." Working Paper, June 30, 2026. <https://ramp.com/data/ai-jobs-impact>
- 22 Michael Cembalest (Chairman of Market and Investment Strategy, J.P. Morgan Asset & Wealth Management). "Eye on the Market, Outlook 2026: Smothering Heights", January 2026: "65%-75% of S&P 500 returns, profits and capital spending since the launch of ChatGPT in 2022 have been derived from 42 companies linked to generative AI" <https://am.jpmorgan.com/content/dam/jpm-am-aem/global/en/insights/eye-on-the-market/smothering-heights-amv.pdf>.
- 23 Damisola Sulaiman: "Starling Bank to cut 130 jobs amid AI push." Sifted, July 3, 2026 (first report: Financial Times, July 3, 2026); financial figures: Starling Bank, Annual Report 2026. <https://sifted.eu/articles/starling-bank-job-cuts>
- 24 Amy Coleman: "The latest in our company transformation." Microsoft Official Blog, July 6, 2026; additionally Todd Bishop: "Microsoft cuts 4,800 jobs, about 2% globally, revamps salesforce and launches massive Xbox overhaul." GeekWire, July 6, 2026. <https://blogs.microsoft.com/blog/2026/07/06/the-latest-in-our-company-transformation/>
- 25 "SAP Cuts Back Hiring, Travel to Fund 'Significant' AI Push." Bloomberg, July 2, 2026; confirmed by The Register, July 2, 2026, and Handelsblatt, July 2026. <https://www.bloomberg.com/news/articles/2026-07-02/sap-restricts-hiring-travel-to-fund-significant-ai-push>

- 26 Katie Paul and Jeff Horwitz: "Mark Zuckerberg had a bold plan to replace Meta staff with AI. Here's how it imploded." Reuters Special Report, August 26, 2026, based on internal planning documents, posts and recordings as well as more than twenty interviewees; Meta confirmed to Reuters the existence and two wave structure of Project OT and the scenario of shrinking individual teams by up to sixty percent, but denied the intention to cut sixty percent of the total workforce, and did not name the units affected. Ten percent reduction of May 20, 2026 (around 8,000 employees) independently documented via NPR, Al Jazeera and CNBC of the same day; headcount 75,472 as of June 30, 2026 and capex guidance of 130 to 145 billion dollars per Meta's Q2 2026 release of July 29, 2026 (SEC Exhibit 99.1); essay Mark Zuckerberg, "The Future is for Everyone", meta.com/thefutureisforeveryone, August 10, 2026. German language coverage: Alexander Amon, "Entlassungswelle eingebremst: Meta scheitert mit KI-Umbau der eigenen Firma", Der Standard, September 1, 2026, <https://www.derstandard.at/story/3000000337646/>
- 27 "Companies Warned AI Would Wipe Out Jobs. Now They're Not So Sure." The Wall Street Journal, July 6, 2026; job cut figures additionally via Challenger, Gray & Christmas: "Challenger Report: June Layoffs Cool to 45,849, Down 53% From May; AI Leads Reasons for Fourth Consecutive Month", July 2026, as well as EY-Parthenon CEO Survey (46 versus 20 percent, documented only by secondary sources). <https://www.wsj.com/tech/ai/ai-workers-tech-ceos-job-losses-afc71e15>
- 28 "How tech workers are feeling in 2026: a workforce splitting in two". Noam Segal and Lenny Rachitsky, Lenny's Newsletter, July 7, 2026. <https://www.lennysnewsletter.com/p/how-tech-workers-are-feeling-in-2026>.
- 29 Gary Hamel and Michele Zanini: "Closing the AI Hype Gap", working paper, June 2026, 13 pages (verified secondary sources within it: Daron Acemoglu, "The Simple Macroeconomics of AI", Economic Policy 40(121), January 2025, TFP gain of 0.53 to 0.66 percent over ten years; Penn Wharton Budget Model, September 8, 2025, AI contribution to TFP growth in 2025 at 0.01 percentage points). <https://garyhamel.com/wp-content/uploads/2026/06/Closing-the-AI-Hype-Gap.pdf>
- 30 Julie Ray: "Americans Cool Toward AI. After Two Years of Improving Attitudes Toward AI, Americans Turn Skeptical". Gallup, July 28, 2026. <https://news.gallup.com/poll/712751/americans-cool-toward.aspx>.
- 31 The AI Disruption: From Doomsday Destruction to Do-Nothing Bots! Aswath Damodaran, Musings on Markets (Blogspot) and Substack, March 2026. <https://aswathdamodaran.blogspot.com/2026/03/ai-scenarios-from-economic-doomsday-to.html>
- 32 Dany Bahar and Greg C. Wright: "New evidence on data center employment effects" (based on the working paper "Data Centers and Local Labor Markets"). Brookings Institution, May 4, 2026. <https://www.brookings.edu/articles/new-evidence-on-data-center-employment-effects/>
- 33 Patrick Franke and Adrian Keppler: "KI: Wer hat Angst vorm digitalen Mann?". Series "Im Fokus: KI", Helaba Research, Landesbank Hessen-Thüringen Girozentrale, August 18, 2026, 6 pages; edited by Dr. Stefan Mitropoulos, responsible Dr. Gertrud Rosa Traud, Chief Economist.
- 34 Andrea Nahles, Chair of the Executive Board of the Bundesagentur für Arbeit, in an interview with Deutschlandfunk: "BA-Vorstandsvorsitzende Nahles: Noch 162.000 freie Ausbildungsstellen für dieses Jahr", published August 15, 2026. <https://www.deutschlandfunk.de/ba-vorstandsvorsitzende-nahles-noch-162-000-freie-ausbildungsstellen-fuer-dieses-jahr-102.html>. The part on the appeal proceedings is reported via dts Nachrichtenagentur of August 16, 2026, among others at HASEPOST, <https://hasepost.de/mehr-ki-widersprueche-bei-arbeitsagentur-743989/>, and in the Oldenburger Onlinezeitung, <https://www.oldenburger-onlinezeitung.de/nachrichten/arbeitsagentur-erfasst-mehr-widersprueche-wegen-ki-223513.html>, as well as in Die Zeit, Arbeit section, August 2026.
- 35 Bundesagentur für Arbeit: "Bürgergeld: Widersprüche und Klagen im Jahr 2025", press release No. 3, Nuremberg, January 13, 2026. <https://www.arbeitsagentur.de/vor-ort/saarland/presse/2026-3-buergergeld-widersprueche-und-klagen-im-jahr-2025>.
- 36 Skills in the AI age. OECD Artificial Intelligence Papers, No. 60, written by Elif Köksal-Oudot under the direction of Molly Leshner, OECD Publishing, Paris, published July 3, 2026. <https://doi.org/10.1787/972bd15e-en>
- 37 U.S. Bureau of Labor Statistics: "Artificial Intelligence (AI) exposure categories", Employment Projections Program, published August 27, 2026. <https://www.bls.gov/emp/publications/ai-exposure-categories.htm>, together with the data table "AI exposure categories and 2025-35 employment projections", <https://www.bls.gov/emp/ind-occ-matrix/ai-exposure-categories.xlsx>, accessed August 28, 2026. Framing release: U.S. Bureau of Labor Statistics, "Employment Projections 2025-2035", News Release, August 27, 2026, <https://www.bls.gov/news.release/ecopro.nr0.htm>.
- 38 AI Economic Indicators: June 2026 Update. Research Note No. 1, Stanford Digital Economy Lab (director Erik Brynjolfsson), June 10, 2026. https://digitaleconomy.stanford.edu/app/uploads/2026/06/AIEL_RN01_Jun26.pdf; platform indicators.stanford.edu
- 39 Anders Humlum and Emilie Vestergaard. "Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI". NBER Working Paper No. 33777, first issued May 2025, revised version cited here of March 13, 2026, 130 pages, DOI 10.3386/w33777. <https://www.nber.org/papers/w33777>. The paper previously ran under the title "Large Language Models, Small Labor Market Effects"; older versions under that title are superseded.

- 40 Massenkoff, M. and McCrory, P. (2026). Labor market impacts of AI: A new measure and early evidence. Anthropic (research note), March 5, 2026. <https://www.anthropic.com/research/labor-market-impacts>
- 41 Analysis of the unemployment rate by AI exposure quintiles (2022 to early 2025): top quintile +0.30 percentage points, bottom quintile +0.94 percentage points. Sarah Eckhardt, Nathan Goldschlag. Economic Innovation Group, August 10, 2025. <https://eig.org/ai-and-jobs-the-final-word/>
- 42 2026 Corporate Recruiters Survey Report. Graduate Management Admission Council (GMAC), with EFMD and Career Services and Employer Alliance, survey period January to May 2026. <https://www.gmac.com/market-intelligence-and-research/research-library/employment-outlook/2026-corporate-recruiters-survey-report>
- 43 Looking for the Ladder: Is AI Impacting Entry-Level Jobs? Zanna Iscenko, Fabien Curto Millet. Economic Innovation Group, The American Worker Project, January 2026 (published alongside on January 14, 2026 on Agglomerations). eig.org/wp-content/uploads/2026/01/TAWP-Iscenko-Millet.pdf
- 44 The Broken Ladder: AI, Remote Work, and Early-Career Hiring. Peter John Lambert, Yannick Schindler. Working Paper, London School of Economics, May 18, 2026. SSRN 6787638, DOI 10.2139/ssrn.6787638, papers.ssrn.com/sol3/papers.cfm?abstract_id=6787638
- 45 Remote Work Leaves Younger Workers Sidelined. Natalia Emanuel, Emma Harrington, Amanda Pallais. Federal Reserve Bank of New York, Liberty Street Economics, June 1, 2026, DOI 10.59576/lse.20260601. libertystreeteconomics.newyorkfed.org/2026/06/remote-work-leaves-younger-workers-sidelined/
- 46 Magnus Lodefalk, Lydia Löthman, Michael Koch, Erik Engberg: "Same Storm, Different Boats: Generative AI and the Age Gradient in Hiring". Working Paper, Örebro University School of Business, ORU Working Paper 2/2026 (Economics), also Ratio Working Paper 388, version v2 of March 16, 2026, review status: revise and resubmit at Economics Letters. <https://www.oru.se/globalassets/oru-sv/institutioner/hh/workingpapers/workingpapers2026/wp-2-2026.pdf>; the most current version in each case at <https://www.ai-econlab.com/papers>, record at <https://swopec.hhs.se/ratio/abs/ratioi0388.htm>.
- 47 Daten- & KI-Skills im Arbeitsmarkt 2026, Österreich. Aloisius Caraet, Simona Hübl, Maximilian Fischer. Nejo and data:unplugged, data as of May 1, 2026, published June 24, 2026. Coverage: Hannah Fasching, "Studie zu KI-Arbeitsmarkt, Nur knapp 8 Prozent Junior-Ausschreibungen", Brutkasten, June 25, 2026, brutkasten.com/artikel/studie-zu-ki-arbeitsmarkt-nur-knapp-8-prozent-junior-ausschreibungen
- 48 Hellwagner, Timon, and Enzo Weber. Generation Z so häufig im Arbeitsmarkt wie junge Leute seit Jahrzehnten nicht. IAB-Forum, August 10, 2026. <https://iab-forum.de/generation-z-so-haeufig-im-arbeitsmarkt-wie-junge-leute-seit-jahrzehnten-nicht/>. Data basis Mikrozensus and statistics of the Bundesagentur für Arbeit. The breakdown by students and non students for the period 2015 to 2023 comes from the earlier IAB press release "Generation Z: Erwerbsbeteiligung der 20- bis 24-Jährigen klettert auf den höchsten Stand seit Jahrzehnten" of February 17, 2025, <https://iab.de/presseinfo/generation-z-erwerbsbeteiligung-der-20-bis-24-jaehrigen-klettert-auf-den-hoechsten-stand-seit-jahrzehnten/>.
- 49 ifo Institut. "Unternehmen erwarten eher sinkende Löhne durch Einsatz von Künstlicher Intelligenz". ifo Konjunkturumfrage, release of August 12, 2026. <https://www.ifo.de/fakten/2026-08-12/unternehmen-erwarten-eher-sinkende-loehne-durch-einsatz-von-kuenstlicher>
- 50 Christian Schumacher, Can Tihanyi, Andreas Schumacher: "State of AI Workforce · Austria · 2026". Self published research report, 70 pages, ISBN 978-3-200-11326-8, May 2026, version 1.0, together with a practitioner version (34 pages). www.stateofai.at, PDFs at stateofai.at/downloads/State_of_AI_Austria_2026.pdf and stateofai.at/downloads/practitioner-report.pdf. Authors: Christian Schumacher (Institute for Strategic Management, WU Wien), Can Tihanyi (Institute for International Business, WU Wien), Andreas Schumacher (Chief Innovation Officer, VTU Group). Reason for inclusion: Michael Köttirsch, "Wie Österreichs KI-Potenzial am Personal scheitert", Die Presse, August 14, 2026, <https://www.diepresse.com/41173110/wie-oesterreichs-ki-potenzial-am-personal-scheitert>.
- 51 "AI-Native Firms". Hyunjin Kim and Rembrand Koning, Harvard Business School Working Paper No. 26-090, June 9, 2026. https://www.hbs.edu/ris/Publication%20Files/26-090_96f92aa0-37d9-4789-beaa-5c0cb87a4032.pdf (in parallel on SSRN: papers.ssrn.com/sol3/papers.cfm?abstract_id=6905079).
- 52 "AEO Job Openings in 2026: Salary Data, Hiring Trends, and Market Insights." Kaleigh Moore, kaleighmoore.com, July 6, 2026. <https://www.kaleighmoore.com/blog/2026/7/6/ai-search-jobs>
- 53 "Gartner Predicts 60% of Organizations Will Adopt Smaller Software Engineering Teams by 2029." Gartner Newsroom, press release (analyst Aliyah Camacho), Stamford, July 7, 2026. <https://www.gartner.com/en/newsroom/press-releases/2026-07-07-gartner-predicts-60-percent-of-organizations-will-adopt-smaller-software-engineering-teams-by-2029>

- 54 Ivan Yotzov, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Steven J. Davis et al.: "Firm Data on AI." NBER Working Paper No. 34836, National Bureau of Economic Research, February 2026 (revised version March 2026); data basis includes the Survey of Business Uncertainty of the Federal Reserve Bank of Atlanta. <https://www.nber.org/papers/w34836>
- 55 David Strömberg (Stockholm University), Victor Lei and Yanhui Wu (University of Hong Kong): "The Generative AI Learning Penalty. Evidence from Chinese Secondary Education". CEPR Discussion Paper DP21577, June 2, 2026, mirrored on SSRN, DOI 10.2139/ssrn.6977138, not peer reviewed.
- 56 "When Everyone Uses AI, Companies Risk Losing Critical Skills." Sagar Goel, David Martin, Charikleia Kaffe. BCG / BCG Henderson Institute, June 17, 2026. <https://www.bcg.com/publications/2026/when-everyone-uses-ai-companies-risk-critical-skills>
- 57 Three sources: (1) Anam Butt, Janine Berg, "Gen AI, occupational segregation and gender equality in the world of work", ILO Research Brief, March 5, 2026, DOI 10.54394/00033798, <https://www.ilo.org/publications/gen-ai-occupational-segregation-and-gender-equality-world-work>; (2) Mauro Cazzaniga et al., "Gen-AI: Artificial Intelligence and the Future of Work", IMF Staff Discussion Note SDN/2024/001, January 2024, <https://www.imf.org/-/media/files/publications/sdn/2024/english/sdnea2024001.pdf>; (3) "Algorithm and Eve: How AI will impact women at work", OECD Policy Brief, December 2024, DOI 10.1787/a1603510-en, https://www.oecd.org/en/publications/2024/12/algorithm-and-eve_0e889c45.html
- 58 Virginia Sondergeld, Katharina Wrohlich, Julia Redelings: "Transformationspotenzial durch KI betrifft ebenso Männer wie Frauen auf dem Arbeitsmarkt". DIW Wochenbericht No. 14+15/2026, volume 93, Deutsches Institut für Wirtschaftsforschung, Berlin, April 1, 2026, pp. 235 to 241, DOI 10.18723/diw_wb:2026-14-1. First author Sondergeld is an economist at the Indeed Hiring Lab and a guest researcher in DIW's Gender Economics research group; the index used comes from Indeed.
- 59 Aditya Challapally, Chris Pease, Ramesh Raskar and Pradyumna Chari. "The GenAI Divide: State of AI in Business 2025". MIT NANDA (Project NANDA, MIT Media Lab), version v0.1, cover date July 2025, 26 pages, widely received from August 2025 onward. The original MIT address nanda.media.mit.edu/ai_report_2025.pdf is no longer available and redirects to the group page; only mirror copies are citable, no DOI. The report describes itself as "Preliminary Findings" without peer review.
- 60 Alex Singla, Alexander Sukharevsky, Lareina Yee and Michael Chui, with Bryce Hall and Tara Balakrishnan. "The state of AI in 2025: Agents, innovation, and transformation". QuantumBlack, AI by McKinsey / McKinsey & Company, November 5, 2025. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>.
- 61 Noy, Shakked and Whitney Zhang. Experimental evidence on the productivity effects of generative artificial intelligence. Science, volume 381, issue 6654, July 13, 2023, pp. 187-192. DOI: 10.1126/science.adh2586.
- 62 Appel, R., Massenkoff, M., McCrory, P. et al. (2026). Anthropic Economic Index report: Economic primitives. Anthropic, January 15, 2026. <https://www.anthropic.com/research/anthropic-economic-index-january-2026-report>
- 63 Iscenko, Z., Strand, S., Manyika, J. et al. AI & Economy ATLAS v1.0: Mapping Gemini Usage in the Economy. Google and Google DeepMind, July 23, 2026. <https://ai.google/static/documents/GoogleATLASv1.pdf>
- 64 Institut für Arbeitsmarkt- und Berufsforschung (IAB) der Bundesagentur für Arbeit, Nuremberg. Substituierbarkeitspotenziale der Berufe 2022, IAB-Kurzbericht 05/2024 (time series 2013-2022), as well as the online tool Job-Futuromat. <https://doku.iab.de/kurzber/2024/kb2024-05.pdf> ; <https://job-futuromat.iab.de/> Additionally the interview with Britta Matthes and Katharina Grienberger, conducted by Christiane Keitel: "Es geht nicht darum, was KI uns wegnehmen könnte, sondern welche Chancen entstehen". IAB-Forum, May 11, 2026, <https://iab-forum.de/es-geht-nicht-darum-was-ki-uns-wegnehmen-koennte-sondern-welche-chancen-entstehen/>. Independently documented by the Verordnung über die Berufsausbildung zum Steuerfachangestellten und zur Steuerfachangestellten of August 3, 2022, Bundesgesetzblatt 2022 Part I Number 30, in force since August 1, 2023, which replaces the regulation of May 9, 1996. The tool itself was rebuilt in January 2026.
- 65 IW Consult GmbH on behalf of eco, Verband der Internetwirtschaft e.V. Der Deutschland-Case: Wie KI das industrielle Geschäftsmodell erneuert. Published July 7, 2026. <https://www.eco.de/der-deutschland-case-wie-ki-das-industrielle-geschaeftsmodell-erneuert/>
- 66 Joachim Nagel: "Already Here, Not Yet Everywhere: Shaping the Economic Impact of Artificial Intelligence", keynote at the International Economic Symposium, Rome, April 21, 2026, Deutsche Bundesbank. <https://www.bundesbank.de/en/press/speeches/already-here-not-yet-everywhere-shaping-the-economic-impact-of-artificial-intelligence-994048>
- 67 David, P. A. (1990). The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox. The American Economic Review, vol. 80, no. 2, May 1990, pp. 355-361. <https://www.jstor.org/stable/2006600>

- 68 "From GenAI to Agentic AI. Operating-model shift in financial services". Wavestone and Institut für Wirtschaftsinformatik der Universität St. Gallen (Eugen Ensinger, Luc-Phillip Heide, Jan Marco Leimeister, Mahei Manhai Li, Teresa Grauer), July 9, 2026. <https://www.wavestone.com/de/insight/agentic-ai-in-financial-services-2/>.
- 69 The Department of State Generative AI Playbook. A practical guide for U.S. government agencies developing and scaling enterprise AI solutions. U.S. Department of State, Bureau of Diplomatic Technology, Center for Analytics, Washington, July 2026, 39 pages. https://www.state.gov/wp-content/uploads/2026/07/DOS-Generative-AI-Playbook_July-2026.pdf
- 70 Notion Research. The Great Renovation: The State of Global AI Transformation. Notion Labs, 2026 (field period March 19 to May 19, 2026, survey partner Qualtrics). <https://www.notion.com/resources/inside-the-ai-transformation>
- 71 G-P (Globalization Partners). AI at Work: The 2026 Reality Check. Fieldwork by Wakefield Research (March 6 to 16, 2026), published May 12, 2026. <https://www.globalization-partners.com/resources/report-ai-at-work-2026/>
- 72 Global AI Pulse Q2 2026. From deployment to value realization. KPMG International, June 2026. <https://assets.kpmg.com/content/dam/kpmgsites/xx/pdf/2026/06/global-ai-pulse-q2.pdf>
- 73 "Uber burned through its entire 2026 AI budget in four months. Now its COO is questioning whether it's worth it." Fortune, May 26, 2026. <https://fortune.com/2026/05/26/uber-coo-ai-spending-tokens-claude-code/>
- 74 Uber Technologies, Inc., Q2 2026 earnings call of August 5, 2026, statements by chief financial officer Balaji Krishnamurthy on cost per token, adoption of the AI coding tools and hiring pace; transcript via Investing.com, confirmed by Larry Dignan: "Uber controls token spend, outlines AI use cases." Constellation Research, August 5, 2026. <https://www.constellationr.com/insights/news/uber-controls-token-spend-outlines-ai-use-cases>
- 75 Praveen Neppalli Naga (Chief Technology Officer, Uber), post on X of August 5, 2026 on the four efficiency levers and on the end of the so called tokenmaxxing era. <https://x.com/praveenTweets>
- 76 Top SaaS Vendors on Ramp (July 2026). Ara Kharazian, Ramp Economics Lab, July 15, 2026. <https://ramp.com/data/top-saas-vendors-on-ramp-july-2026>
- 77 Direction de la recherche, des études, de l'évaluation et des statistiques (DREES): "Fiche 12, L'imagerie médicale". In: Les dépenses de santé en 2024. Résultats des comptes de la santé, Édition 2025, series Panoramas de la DREES, pages 78 to 80, published September 30, 2025. <https://www.drees.solidarites-sante.gouv.fr/sites/default/files/2025-09/CNS%20-%20Fiche%2012%20-%20L%27imagerie%20m%C3%A9dicale.pdf>.
- 78 Katja Schwabe, Denis Ferizaj, Sandra Neumann, Sandra Strube-Lahmann and Nils Lahmann. "Time Savings Through an AI Speech Assistant for Nursing Documentation: Pre-Post Time-Motion Study in German Long-Term Care". Journal of Medical Internet Research 2026;28:e86078, published April 8, 2026. DOI 10.2196/86078. <https://www.jmir.org/2026/1/e86078/>.
- 79 Microsoft WorkLab, Breaking down the infinite workday (Special Report, June 17, 2025, telemetry data and Edelman Data x Intelligence survey); as well as Mark, G., Gudith, D., Klocke, U. (2008). The Cost of Interrupted Work: More Speed and Stress. ACM CHI 2008.
- 80 Hamsa Bastani (Wharton School), Osbert Bastani (Department of Computer and Information Science, School of Engineering and Applied Science), Alp Sungu, Haosen Ge, Özge Kabakçı and Rei Mariman, University of Pennsylvania: "Generative AI without guardrails can harm learning. Evidence from high school mathematics". Proceedings of the National Academy of Sciences 122(26), e2422633122, June 25, 2025, DOI 10.1073/pnas.2422633122, peer reviewed. <https://www.pnas.org/doi/10.1073/pnas.2422633122> The correction in PNAS 122(34), DOI 10.1073/pnas.2518204122 of August 20, 2025, concerns only an affiliation detail and no numerical value.
- 81 Taha Choukhmane, Tim de Silva, Weidong Lin and Matthew Akuzawa. "AI Financial Advice: Supply, Demand, and Life Cycle Implications". arXiv:2608.01607v1, August 3, 2026, 92 pages. Working paper without peer review, run in parallel as SSRN Working Paper No. 6446286. <https://arxiv.org/abs/2608.01607>.
- 82 Ibid. (2).
- 83 Bret Taylor and Clay Bavor. "Better customer experiences, built on Sierra". Company blog of Sierra, May 4, 2026: "We're raising \$950 million from new and existing investors, led by Tiger Global and GV, at a valuation of over \$15 billion." <https://sierra.ai/blog/better-customer-experiences-built-on-sierra>, verbatim from Taylor on X.
- 84 Measuring AI Ability to Complete Long Tasks. Kwa, T. et al., METR, arXiv:2503.14499, March 2025. <https://arxiv.org/pdf/2503.14499>
- 85 "Can AI agents conduct open-ended AI research? Early evidence from two case studies". Peter Kirgis, Sayash Kapoor, Arvind Narayanan and 21 further authors, CRUX project (Princeton University), arXiv preprint arXiv:2607.27191, submitted July 29, 2026. <https://arxiv.org/abs/2607.27191>.
- 86 Renting is stressful. Millions of renter conversations tell us why. a16z New Media (Andreessen Horowitz), July 23, 2026. <https://www.a16z.news/p/renting-is-stressful-millions-of>

- 87 Investments of the large cloud vendors for 2026. Key figures: Financial Times, April 30, 2026; Morgan Stanley, May 3, 2026; CNBC, February 6, 2026; Bank for International Settlements, Annual Report 2026.
- 88 Arman Shehabi, Alex Newkirk, Sarah J. Smith, Alex Hubbard, Nuoa Lei, Md Abu Bakar Siddik, Billie Holecek, Jonathan Koomey, Eric Masanet and Dale Sartor. "2024 United States Data Center Energy Usage Report". Lawrence Berkeley National Laboratory on behalf of the U.S. Department of Energy, LBNL-2001637, December 2024, DOI 10.71468/P1WC7Q. https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report_1.pdf. Updated by the "United States Data Center Energy Usage Report: 2025 Update", LBNL, June 2026.
- 89 Bank for International Settlements. "Annual Economic Report 2026", chapter I "Progress and peril", June 28, 2026: "The scale and pace of the current AI investment boom accompanied by expectations of large productivity payoffs bear resemblance to these precedents, highlighting potential downside risks in the near term." <https://www.bis.org/publ/arpdf/ar2026e1.htm>. For the U.S. Treasury Department there is no public primary source; it is an internal, so far unpublished draft analysis by career staff to the Treasury Secretary and supervisors, which NOTUS reported exclusively on July 6, 2026. <https://www.notus.org/economy/treasury-internal-report-warning-dangers-ai-bubble>.
- 90 The Nature of the Firm. Ronald H. Coase, *Economica*, New Series, Vol. 4, No. 16, 1937, pp. 386-405.
- 91 Korn Ferry: Workforce 2025: Power Shifts, April 7, 2025, kornferry.com/insights/featured-topics/workforce-management-articles/workforce-planning-insights; Fast Company, Ashley Herd: Everyone wants to kill the middle manager role. The data says don't do it, June 1, 2026.
- 92 The agentic organization: Contours of the next paradigm for the AI era. Alexander Sukharevsky, Alexis Krivkovich, Arne Gast, Arsen Storozhev, Dana Maor, Deepak Mahadevan, Lari Hämäläinen, Sandra Durth, McKinsey & Company, September 26, 2025. mckinsey.com/capabilities/people-and-organizational-performance/our-insights/the-agentic-organization-contours-of-the-next-paradigm-for-the-ai-era
- 93 Warp CEO Zach Lloyd on why software factories are the next phase of coding. Richard MacManus (interview with Zach Lloyd), Latent Space, July 1, 2026. latent.space
- 94 From anxiety to advantage: A marketing organization that thrives with AI. Jacob Ader, Julien Boudet, Kelsey Robinson, Heather Gouinlock, Jim Lecinski, McKinsey & Company, Growth, Marketing & Sales Practice, June 22, 2026. <https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/from-anxiety-to-advantage-a-marketing-organization-that-thrives-with-ai>
- 95 Human-Centered Large Language Models. Reflections and New Directions. Dora Zhao, Caleb Ziems, Ahmad Rushdi, James Landay, Diyi Yang, Stanford University, Institute for Human-Centered Artificial Intelligence (HAI), Industry Report, June 2026. hai.stanford.edu
- 96 Official labor market statistics of the EU and the DACH countries: Eurostat EU Labour Force Survey (Statistical Office of the European Union), Statistisches Bundesamt (Destatis, Germany), Statistik Austria, Bundesamt für Statistik (BFS, Switzerland) as well as Bundesagentur für Arbeit (Germany). <https://ec.europa.eu/eurostat> ; <https://www.destatis.de/> ; <https://www.statistik.at/> ; <https://www.bfs.admin.ch/>
- 97 Eurostat. Statistical Office of the European Union, EU Labour Force Survey and ICT usage in enterprises (database table `isoc_eb_ai`); Eurostat News, "20% of EU enterprises used Artificial Intelligence in 2025", December 11, 2025. <https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20251211-2>
- 98 Cedefop, European Centre for the Development of Vocational Training. Skills Intelligence, Skills Forecast 2025 (Technical Report) and Skills-OVATE (Online Vacancy Analysis Tool for Europe). <https://www.cedefop.europa.eu/en/tools/skills-intelligence>
- 99 Stanford Institute for Human-Centered Artificial Intelligence (HAI). "The 2026 AI Index Report", chapter Economy. Stanford University, 2026. <https://hai.stanford.edu/ai-index/2026-ai-index-report/economy>. Stanford HAI, "The 2025 AI Index Report", <https://hai.stanford.edu/ai-index/2025-ai-index-report>.
- 100 Rodríguez-Pose, Andrés and Federico Bartalucci. A Union under siege, or a Union besieging itself? The EU's new budget and the future of European integration. *Global Challenges & Regional Science*, 7 (2026), 100037, online June 19, 2026. DOI: 10.1016/j.gcrs.2026.100037.
- 101 AI-Econ Lab (2026): AIEL Monitor, modules Exposure, Demand and Adoption. Örebro University and Ratio. <https://ai-econlab.com/monitor/>, accessed August 14, 2026.
- 102 Bonin, Holger, Martin Ertl, Ines Fortin, Daniela Grozea-Helmenstein, Jaroslava Hlouskova, Helmut Hofer, Sebastian P. Koch, Robert M. Kunst, Michael Reiter, Felix Schröter and Klaus Weyerstraß. Export stützt verhaltenes Wachstum. Prognose der österreichischen Wirtschaft 2026 bis 2030. *Wirtschaftsprognose / Economic Forecast No. 143*, Institut für Höhere Studien (IHS), Vienna, July 2026 (embargo July 16, 2026, 10:00 a.m.). <https://www.ihs.ac.at/de/datencenter/konjunkturprognose/>

- 103 The Stepstone Group. "Stepstone Analyse: Weniger Einstiegsjobs, längere Bewerbungsprozesse". Press release, August 12, 2025. https://www.eqs-news.com/news/media/stepstone-analyse-weniger-einstiegsjobs-laengere-bewerbungsprozesse/0150ee69-c13f-46d2-bee7-06734cdb295c_de.
- 104 Verein Deutscher Ingenieure (VDI) and Institut der deutschen Wirtschaft (IW Köln), Axel Plünnecke: "Ingenieurmonitor 2025/IV: Der regionale Arbeitsmarkt in den Ingenieurberufen. Sonderteil: Ausblick bis 2040", published August 3, 2026, reporting period fourth quarter 2025, data basis special analyses by the Bundesagentur für Arbeit. <https://www.iwkoeln.de/studien/axel-pluenecke-der-regionale-arbeitsmarkt-in-den-ingenieurberufen-sonderteil-ausblick-bis-2040.html>. Secondary sources: heise online, "IT-Jobs: Weniger Stellenangebote, Arbeitslosigkeit steigt", August 4, 2026, <https://www.heise.de/news/IT-Jobs-Weniger-Stellenangebote-Arbeitslosigkeit-steigt-11398441.html>; ingenieur.de, August 4, 2026.
- 105 Bundesagentur für Arbeit: "Arbeitsmarkt in der Informations- und Kommunikationstechnik (IKT) im Spannungsfeld: Konjunkturelle Schwäche trifft auf strukturellen Wandel". Press release No. 24, Nuremberg, July 3, 2026, based on the report "Arbeitsmarkt IT-Fachleute" of the BA-Statistik. <https://www.arbeitsagentur.de/presse/2026-24-arbeitsmarkt-in-der-informations-und-kommunikationstechnik-ikt-im-spannungsfeld-konjunkturelle-schwaechе-trifft-auf-strukturellen-wandel>. Secondary source: heise online, August 4, 2026.
- 106 Jeremias Kläui and Michael Siegenthaler. "KI und der Schweizer Arbeitsmarkt. Erste Evidenz zu Auswirkungen auf Arbeitslosigkeit und Stellenausschreibungen". KOF Studien No. 186, KOF Konjunkturforschungsstelle der ETH Zürich, October 2025, accompanying release of October 27, 2025. https://ethz.ch/content/dam/ethz/special-interest/dual/kof-dam/documents/newsletter/KOF_Studie_KI_Schweizer_Arbeitsmarkt.pdf.
- 107 Annalisa Ferrando, Sofia Lamboglia, Judit Rariga and Mathis Schmidt. "Adoption and investment in AI across the euro area: Insights from harmonised firm-level data". ECB Occasional Paper No. 395, July 28, 2026, DOI 10.2866/6856108. <https://www.ecb.europa.eu/pub/pdf/scpops/ecb.op395.en.pdf>. Additionally ECB, Economic Bulletin 2/2026, box of March 31, 2026, and ECB Blog by Laura Lebastard and David Sondermann of March 4, 2026.
- 108 Erik Engberg, Holger Görg, Magnus Lodefalk, Farrukh Javed, Martin Längkvist, Natália Monteiro, Hildegunn Kyvik Nordås, Giuseppe Pulito, Sarah Schroeder and Aili Tang. "AI Unboxed and Jobs: A Novel Measure and Firm-Level Evidence from Three Countries". IZA Discussion Paper No. 16717, Institute of Labor Economics, Bonn, January 2024. <https://docs.iza.org/dp16717.pdf>. Foundational paper of the Dynamic AI Occupational Exposure Index (DAIOE) of the AI-Econ Lab at Örebro University, index data publicly available at <https://www.ai-econlab.com/ai-exposure-daioe>. Working paper without peer review.
- 109 Marc Brupbacher, Olaf König, Patrick Vögeli. "Wie stark lässt sich Ihr Job automatisieren?" (series "400 Berufe im KI-Check"). Tages-Anzeiger, August 2, 2026. <https://www.tagesanzeiger.ch/ki-im-job-welche-berufe-in-der-schweiz-am-staerksten-gefaehrdet-sind-698429113886>. Subscription article.
- 110 OECD. Indicators of Employment Protection, dataset DSD_EPL (OECD.ELS.JAI), country values through reference year 2025, OECD average through 2024, scale 0 to 6. Methodology: OECD (2020), "Recent trends in employment protection legislation", in: OECD Employment Outlook 2020. [https://data-explorer.oecd.org/vis?df\[ds\]=DisseminateFinalDMZ&df\[id\]=DSD_EPL@DF_EPL&df\[ag\]=OECD.ELS.JAI](https://data-explorer.oecd.org/vis?df[ds]=DisseminateFinalDMZ&df[id]=DSD_EPL@DF_EPL&df[ag]=OECD.ELS.JAI).
- 111 Germany: Bundesagentur für Arbeit, Geschäftsbericht 2020, Nuremberg: "In April 2020, the month with the highest take-up, companies received short-time work benefits for 6.00 million people." https://www.arbeitsagentur.de/datei/geschaeftsbericht-2020_ba032090.pdf. Austria: AMS Österreich, Spezialthema zum Arbeitsmarkt 08/2020, at the end of April 2020 1,010,267 people were on short-time work, 29 percent of those actively employed; in its report "COVID-19-Kurzarbeit" (series BUND 2022/7, February 2022) the Rechnungshof puts the actually settled peak in April 2020 at 1,043,152 people.
- 112 McKinsey & Company. Executive survey from early 2018 on retraining versus new hiring (94 percent of European versus 62 percent of American executives).
- 113 Ingka Group: "AI and Remote Selling bring IKEA design expertise to the many", newsroom release, June 29, 2023, <https://www.ingka.com/newsroom/ai-and-remote-selling-bring-ikea-design-expertise-to-the-many/>. Updated by Fortune: "Inside Ikea's big bet on humans in the age of AI", July 30, 2026, <https://fortune.com/2026/07/30/ikea-ai-workforce-reskilling-jobs-billie-chatbot-global-500/>. Job cuts in 2026 per Personnel Today: "Ikea retailer to cut 800 office jobs", March 20, 2026, <https://www.personneltoday.com/hr/ikea-retailer-to-cut-800-office-jobs/>, and Retail Dive: "Ikea franchisor to cut 850 jobs", May 2026, <https://www.retaildive.com/news/inter-ikea-group-layoffs-cut-850-jobs/820614/>.
- 114 Shapiro, Dan. 0%, 50%, or 200%. Blog post, July 28, 2026. <https://www.danshapiro.com/blog/2026/07/0-50-or-200/>
- 115 Institut für Arbeitsmarkt- und Berufsforschung (IAB) der Bundesagentur für Arbeit (headed by Enzo Weber): IAB press release "IAB-Prognose für 2025/2026: Arbeitsmarkt profitiert von Fiskalpaketen, wird aber durch den demografischen Wandel gebremst", September 24, 2025, <https://iab.de/presseinfo/iab-prognose-fuer-2025-2026-arbeitsmarkt-profitiert-von-fiskalpaketen-wird-aber-durch-den-demografischen-wandel-gebremst/>, confirmed by the IAB spring prediction, reported via dpa-AFX on March 24, 2026 ("IAB: Deutsches Erwerbspersonenpotenzial sinkt 2026 erstmals").

- 116 Markus Hoch and Philipp Kreuzer. "Demografie als Wachstumshemmnis. Arbeitsmarktszenarien für Österreich". Prognos AG, Freiburg, November 2025, communicated on December 10, 2025. https://www.prognos.com/sites/default/files/2025-11/202511_Prognos_Demografie_als_Wachstumshemmnis_Arbeitsmarktszenarien_Oesterreich.pdf.
- 117 Robert Stehrer, Laura Zilian and Marcus Knapp. Study by the Wiener Institut für Internationale Wirtschaftsvergleiche (wiiw) on behalf of the Bundesministerium für Arbeit und Wirtschaft on the Austrian labor shortage, published January 13, 2025. <https://wiiw.ac.at/studie-arbeitskraeftemangel-koennte-oesterreichs-wachstum-mittelfristig-gefaehrden-n-663.html>.
- 118 economiesuisse: Dossier Politik "Demografische Entwicklung: Die Pensionierungswelle reisst grosse Lücken auf", <https://www.economiesuisse.ch/de/dossier-politik/demografische-entwicklung-die-pensionierungswelle-reisst-grosse-luecken-auf>; additionally Bundesamt für Statistik (BFS), Bevölkerungsszenarien / "Generationen auf dem Arbeitsmarkt", <https://www.bfs.admin.ch/bfs/de/home/statistiken/arbeit-erwerb/erwerbstaetigkeit-arbeitszeit/alter-generationen-pensionierung-gesundheit/generationen-arbeitsmarkt.html>.
- 119 Staatssekretariat für Wirtschaft (SECO): "Arbeitskräftemangel in der Schweiz", Grundlagen für die Wirtschaftspolitik / Arbeitsmarktpolitik Working Paper No. 67, April 2026; the report of the Bundesrat based on it, "Ursachen und Auswirkungen von Arbeitskräftemangel", adopted April 22, 2026, <https://www.admin.ch/de/newsb/SO23GEuldEZhLmZ2Q9IYB>.
- 120 U.S. Census Bureau. "2023 National Population Projections", main variant and alternative scenarios, published November 9, 2023, press release CB23-189. <https://www.census.gov/programs-surveys/popproj.html>.
- 121 "Situation Älterer am Arbeitsmarkt", series Blickpunkt Arbeitsmarkt, Statistik der Bundesagentur für Arbeit, Nuremberg, November 2024, reference date June 30, 2023 (shares of those aged 55 to under 65 by economic sector). <https://statistik.arbeitsagentur.de/DE/Statischer-Content/Statistiken/Themen-im-Fokus/Demografie/Generische-Publikationen/Aeltere-amArbeitsmarkt.pdf>. Supplemented at the occupational level by Jurek Tiedemann and Franziska Arndt, "Wertvolle Erfahrung sichern", KOFA-Studie 4/2026, Institut der deutschen Wirtschaft on behalf of the Bundesministerium für Wirtschaft und Energie.
- 122 Petra Gregoritsch and Adila Sahbegovic, "Älterenbeschäftigungsmonitoring 2025", Bundesministerium für Arbeit, Soziales, Gesundheit, Pflege und Konsumentenschutz, March 2026, based on the AMS Data Warehouse and the Dachverband der Sozialversicherungsträger. https://www.sozialministerium.gv.at/dam/jcr:65ded023-98f6-4b27-a98d-27fd0e43f18f/260311_Beschaeftigungsmonitoring_Aelterer_2025.pdf
- 123 Bundesamt für Statistik, Schweizerische Arbeitskräfteerhebung, share of those aged 55 to 64 among the employed by economic sector, time series 1995 to 2025. <https://www.bfs.admin.ch/bfs/de/home/statistiken/arbeit-erwerb/erwerbstaetigkeit-arbeitszeit/merkmale-arbeitskraefte/wirtschaftsabschnitt.html>.
- 124 Demis Hassabis: "A Framework for Frontier AI and the Dawning of a New Age", published July 14, 2026 on demishassabis.substack.com (at the same time as a LinkedIn Pulse article). <https://demishassabis.substack.com/p/a-framework-for-frontier-ai-and-the-dawning-of-a-new-age>
- 125 Future of Life Institute: "AI Safety Index Summer 2026", July 2026. <https://futureoflife.org/ai-safety-index-summer-2026/>
- 126 Compiled from: Stanford HAI, "The 2025 AI Index Report", chapter Research and Development, <https://hai.stanford.edu/ai-index/2025-ai-index-report>; Fortune, "US-China AI race: How Tsinghua University is outpacing the Ivy League", November 19, 2025, <https://fortune.com/2025/11/19/us-china-ai-race-higher-education-tsinghua-university-outpacing-ivy-league-mit-harvard-stanford/>; MacroPolo, "The Global AI Talent Tracker 3.0", <https://macropolo.org/digital-projects/the-global-ai-talent-tracker/>; Euronews, "The AI brain drain: why Europe can't keep the talent it trains", January 2026, <https://www.euronews.com/my-europe/2026/01/29/the-ai-brain-drain-why-europe-cant-keep-the-talent-it-trains>.
- 127 Daniel Kokotajlo, Scott Alexander, Thomas Larsen, Eli Lifland, Romeo Dean: "AI 2027. A research-backed AI scenario forecast", AI Futures Project, April 3, 2025. <https://ai-2027.com/>
- 128 Thomas Larsen, Romeo Dean, Brendan Halstead, Eli Lifland, Ryan Greenblatt, Daniel Kokotajlo (sections revised by Scott Alexander): "AI 2040: Plan A", AI Futures Project, July 9, 2026. <https://ai-2040.com>
- 129 Julia Willemyns: "Are you afraid of the Luddites? On the political economy of AI in China, America, and Europe", Substack essay, July 24, 2026. <https://juliawillemyns.substack.com/p/are-you-afraid-of-the-luddites>
- 130 Anton Korinek, Charles I. Jones, Szymon Sacher, Tess Cotter and Peter McCrory: "Economic Scenarios for Transformative AI". The Anthropic Institute, Working Paper No. 2026-02, September 2026, 57 pages. <https://www-cdn.anthropic.com/files/4zrzovbb/website/cf58f84d46a4a76bf5a5b039ac695fba6b80041c.pdf>, accompanying Scenario Explorer (version 1.0) at <https://www.anthropic.com/institute/econ-scenarios>.
- 131 "Why wages and productivity look set to diverge further". Joel Suss, Financial Times (Free Lunch Newsletter), published July 26, 2026. <https://www.ft.com/content/9741206d-b53f-4149-8cab-5203dd4b70fb>.

- 132** "What Humanity Needs To Flourish In The Next Decade". E. Glen Weyl, James Evans and Chris White, with a 23 member coalition from research, business and civil society, Noema Magazine (Berggruen Institute), July 21, 2026. <https://www.noemamag.com/what-humanity-needs-to-flourish-in-the-next-decade/>.
- 133** "The strategic role of higher education to deliver the Government's Industrial Strategy". University of York, QS Quacquarelli Symonds and Public First, January 2026. <https://www.qs.com/insights/higher-education-impact-governments-industrial-strategy-skills>.
- 134** "The Future Worth Building Is Human". Thinking Machines Lab, company blog, July 10, 2026. thinkingmachines.ai/blog/the-future-worth-building-is-human/.
- 135** AXA Versicherungen AG (commissioned) and research institute Sotomo, Zurich (conducted; Lisa Frisch, Severin Meier, Simon Stückelberger): "AXA KMU-Arbeitsmarktstudie 2026", company survey, July 2026, 62 pages. Online survey of 336 Swiss SMEs with five or more employees, March 11 to 18, 2026, weighted by size and sector. AI chapter pages 38 to 47, figures 22 to 24; all values used were checked against the figures, the year markers in the study's running text deviate from them. Media release on the AI section dated September 15, 2026. PDF 20260709_axa_kmu-arbeitsmarktstudie.pdf via axa.ch.
- 136** Benjamin Verschuere and Angus Cameron: "AI Impact on Jobs". Liminal Capital, August 2026. <https://www.liminal-capital.com/documents/liminal-ai-jobs-impact.pdf>; extended version as "Hiding in the Mean: The Two Margins of AI's Employment Effect", SSRN Working Paper No. 7195359, and "What Economists Are Missing About AI", ProMarket (Stigler Center, University of Chicago Booth School of Business), July 15, 2026. Data basis: 17,536 task descriptions from the O*NET occupational database, Occupational Employment and Wage Statistics and Current Population Survey of the Bureau of Labor Statistics 2015 to 2025, value added data from the Bureau of Economic Analysis. The front page story on it: Talmon Joseph Smith: "The Quiet Way A.I. Is Hitting the Work Force", The New York Times, September 16, 2026, business section. Interests: quantitative macro fund, no peer review, task classification by language models. Not used are the long term estimate of around 20 million positions and the coefficients of the regression model.
- 137** Torsten Slok (chief economist, Apollo Global Management) and Sania Edlich (Princeton University): "AI Lowers Wages But Doesn't Cut Jobs". The Daily Spark, Apollo Global Management, July 30, 2026. <https://www.apollo.com/wealth/insights-news/insights/daily-spark/ai-lowers-wages-but-doesnt-cut-jobs>. Difference in differences across 321 matched occupations, 2015 to 2026, exposure measure is the Anthropic Economic Index, that is observed use rather than theoretical feasibility. Independently reported by Axios, July 31, 2026, and Dataconomy, August 24, 2026. Additionally cited in The New York Times of September 16, 2026, there rounded to 11 percent for the lowest wage quartile. Interests: asset manager, no peer review. Not adopted are the value of 24 percent for service occupations as a small subsample, any transfer to the German speaking region, and any link between the values of this analysis and the earlier calculation across eleven occupations and around 5.8 million employees.
- 138** Talmon Joseph Smith: "The Quiet Way A.I. Is Hitting the Work Force". The New York Times, September 16, 2026, business section. <https://www.nytimes.com/2026/09/16/business/ai-raises-hiring.html>. Used are the juxtaposition of Dario Amodeli's prediction from May 2025 with the official unemployment rate of 4.1 percent in August 2026 and the verbatim quote from Benjamin Verschuere. Full text read via the publisher's gift link, archive excerpt in the sources folder. Not used are the industry figures for the technology sector and financial services and the three company examples, all singly sourced or self reported.
- 139** U.S. Bureau of Labor Statistics: "Labor share at its lowest level, 52.8 percent, in second quarter 2026". The Economics Daily, September 15, 2026. <https://www.bls.gov/opub/ted/2026/labor-share-at-its-lowest-level-52-8-percent-in-second-quarter-2026.htm>; series Nonfarm Business Sector: Labor Share for All Workers (PRS85006173), data from 1947. Official statistics without commercial interests. The figure circulating in the press of a decline of around 10 percentage points since a peak in the late 1990s is not adopted, because the official series shows its peak in the fourth quarter of 1960.
- 140** Virginia Sondergeld: "Berufseinstieg: Stellenangebot sinkt, Ausbildungssystem bietet einen Anker". Indeed Hiring Lab Germany, September 23, 2026. <https://hiringlab.indeed.com/de/blog/2026/09/23/berufseinstieg-stellenangebot-sinkt-ausbildungssystem-bietet-einen-anker/>. Analysis of job postings on Indeed Germany since 2019, reference value is the 2019 annual average; entry-level positions, internships and apprenticeships are delimited by terms in the job title, the AI transformation potential follows the Indeed AI at Work Report 2025 based on American postings from May 2024 to April 2025. Provider data of a commercial job board, not peer-reviewed; what is measured are postings, not hires and not employment.
- 141** Dana Hajek: "Arbeitsstellenrückgang: Wo Einstiegsjobs wegen KI wegfallen". Die Zeit, September 18, 2026. <https://www.zeit.de/arbeit/2026-09/arbeitsstellenrueckgang-berufseinstiegs-kuenstliche-intelligenz-arbeitsmarkt>. Charts based on the Indeed data with reference year 2020; the values cited in the report follow the Hiring Lab publication with reference year 2019. Quotations from Melanie Arntz (Institute for Employment Research) and Virginia Sondergeld (Indeed Hiring Lab).

Version history of the 2026 edition

Version 1.0 September 16, 2026 First publication of the 2026 edition. English edition of the German original published on the same day.

Version 1.1 September 18, 2026 Gender chapter and key finding 5 switched to the Eurostat 2025 statistics on the use of generative AI (previously IAB/D21, survey of July 2024); ninth prediction on the gender gap at work added, the previously deferred prediction on the D21 series dropped; source 6 extended. Newly added as source 135, the Swiss SME labor market study 2026 by AXA and Sotomo: the expectation of job cuts among AI-using SMEs at 12 percent for three years (chapter on the AI argument), trial versus operation and the decline of text tasks after the 2025 peak (chapter on the long rebuild). The source numbers of version 1.0 remain unchanged; new sources are appended.

Version 1.2 September 24, 2026 The assessments from the author's working conversations have been dropped; since this version, the report relies exclusively on verifiable sources. The foreword and the method section have been adjusted accordingly, as have the affected passages in the chapters, in the predictions and in the final chapter. Newly added as sources 136 and 137: the task decomposition by Liminal Capital, according to which career entries of 22 to 25 year olds in substitutable occupations have fallen by about a third since 2021 while layoffs remained unchanged (chapter on the career gap), and the wage finding by Apollo Global Management, wage growth 6.7 percentage points slower in exposed occupations with employment unchanged (final chapter); as sources 138 and 139 the New York Times article of September 16, 2026 and the official labor share from the Bureau of Labor Statistics (52.8 percent in the second quarter of 2026, the lowest since the series began in 1947), which give the Amodei benchmark a first interim status and extend the distribution paragraph in the final chapter. Newly added as sources 140 and 141: the Indeed Hiring Lab analysis of the German entry-level job market since 2019 (entry-level postings 46 percent below 2019, apprenticeships 50 percent above, divergence by AI exposure since spring 2022) together with the Die Zeit article of September 18, 2026 with Melanie Arntz's assessment; author attribution corrected for source 58 (first author at the Indeed Hiring Lab). Method section: the rhythm of the report is now fixed. Three levels, the annual edition, the quarterly version and the analyses on workafterai.org, with the rules that predictions remain unchanged within an edition and a version appears only when a finding has changed or an error needs correcting; version 1.3 is planned for early November and the final version 1.4 for December 31, 2026. As of this version, the publisher is Kuerner Group GmbH, Wels; the author remains Gerhard Kürner. Work DOI, ORCID and the new citation recommendation are in the method section. With this version, the report also appears for the first time as a machine edition, that is as a text edition with a fixed identifier for every chapter, key finding, prediction, external benchmark and source, as data files and as a skill package; the files sit next to the PDF under the work DOI. The source numbers of version 1.1 remain unchanged; new sources are appended.